



**NEW UTILITY APPLICATION REQUEST
TRANSMITTAL UNDER 37 C.F.R. §1.53(b)**

Docket No: G0008/7003



CERTIFICATE OF EXPRESS MAILING

"Express Mail" mailing label number: **EL575396898**

Date of Deposit: **June 6, 2000**

I hereby certify that the following Correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service pursuant to 37 C.F.R. §1.10 on the date indicated above in an envelope addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231.

Jan L. Mellen
Jan L. Mellen

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 20231

This is a request for a filing of a new utility patent application under 37 C.F.R. §1.53(b) for the following:

Inventor(s): Raymond E. Ozzie, Kenneth G. Moore, Ransom L. Richardson And Edward J. Fischer

Title: METHOD AND APPARATUS FOR EFFICIENT MANAGEMENT OF XML DOCUMENTS

Enclosures

123 Pages of Specification, including claims and abstract
14 Sheets of Formal Drawing(s)

- ☒ Declaration and Power of Attorney
- ☒ Assignment(s) of the Invention and Assignment Recordation Cover Sheet
- ☐ Small Entity Statement
- ☐ PTO 1449 and a copy of each cited reference
- ☒ Return Receipt Postcard
- ☐ Other:

Filing Fees

Claims as Filed					
	Claims Filed	Basic Fee Allowance	Number of Extra Claims	Rate	Fees Due
Basic Fee 37 CFR §1.16(a)					\$690.00
Total Claims (37 CFR §1.16(c))	82	- 20 =	62 X	\$18.00 =	\$1,116.00
Independent Claims (37 CFR §1.16(b))	16	- 3 =	13 X	\$78.00 =	\$1,014.00
Reduction by 50% for filing by small entity					\$1,410.00
Assignment Recording Fee					\$ 40.00
Total Filing Fee					<u>\$1,450.00</u>

Payment

- ☒ Check in the amount of the total filing fee is enclosed.
☐ Charge Account No. 02-3038 in the amount of the total filing fee.
A duplicate of this transmittal is attached.

Authorization to Charge Additional Fees

- ☒ The Commissioner is hereby authorized to charge any additional fees incurred under 37 C.F.R. §1.16, §1.17 and §1.18 required by this paper and during the entire pendency of this application to Account No. 02-3038.

Correspondence Address

Please forward all correspondence to
Paul E. Kudirka, Esq. at the address of customer no:

021127



Paul E. Kudirka

Date: 6/6/00

Paul E. Kudirka, Esq. Reg. No. 26,931
KUDIRKA & JOBSE, LLP
Customer Number 021127
Tel: (617) 367-4600 Fax (617) 367-4656

PATENT: UTILITY

Docket No: G0008/7003

Inventors: Raymond E. Ozzie, Kenneth G. Moore, Ransom L. Richardson and
Edward J. Fischer

**METHOD AND APPARATUS FOR EFFICIENT
MANAGEMENT OF XML DOCUMENTS**

METHOD AND APPARATUS FOR EFFICIENT MANAGEMENT OF XML DOCUMENTS

FIELD OF THE INVENTION

5 This invention relates to storage and retrieval of information and, in particular, to storage and retrieval of information encoded in Extended Markup Language (XML).

BACKGROUND OF THE INVENTION

10 Modern computing systems are capable of storing, retrieving and managing large amounts of data. However, while computers are fast and efficient at handling numeric data they are less efficient at manipulating text data and are especially poor at interpreting human-readable text data. Generally, present day computers are unable to understand subtle context information that is necessary to understand and recognize pieces of information that comprise a human-readable text document. Consequently, 15 although they can detect predefined text orderings or pieces, such as words in an undifferentiated text document, they cannot easily locate a particular piece of information where the word or words defining the information have specific meanings. For example, human readers have no difficulty in differentiating the word "will" in the sentence "The attorney will read the text of Mark's will.", but a computer may have great 20 difficulty in distinguishing the two uses and locating only the second such use.

25 Therefore, schemes have been developed in order to assist a computer in interpreting text documents by appropriately coding the document. Many of these schemes identify selected portions of a text document by adding into the document information, called "markup tags", which differentiates different document parts in such a way that a computer can reliably recognize the information. Such schemes are generally called "markup" languages.

30 One of these languages is called SGML (Standard Generalized Markup Language) and is an internationally agreed upon standard for information representation. This language standard grew out of development work on generic coding and mark-up languages, which was carried out in the early 1970s. Various lines

of research merged into a subcommittee of the International Standards Organization called the subcommittee on Text Description and Processing Languages. This subcommittee produced the SGML standard in 1986.

SGML itself is not a mark-up language in that it does not define mark-up tags nor does it provide a markup template for a particular type of document. Instead, SGML denotes a way of describing and developing generalized descriptive markup schemes. These schemes are generalized because the markup is not oriented towards a specific application and descriptive because the markup describes what the text represents, instead of how it should be displayed. SGML is very flexible in that markup schemes written in conformance with the standard allow users to define their own formats for documents, and to handle large and complex documents, and to manage large information repositories.

Recently, another development has changed the general situation. The extraordinary growth of the Internet, and particularly, the World Wide Web, has been driven by the ability it gives authors, or content providers, to easily and cheaply distribute electronic documents to an international audience. SGML contains many optional features that are not needed for Web-based applications and has proven to have a cost/benefit ratio unattractive to current vendors of Web browsers. Consequently, it is not generally used. Instead, most documents on the Web are stored and transmitted in a markup language called the Hypertext Markup Language or HTML.

HTML is a simple markup language based on SGML and it is well suited for hypertext, multimedia, and the display of small and reasonably simple documents that are commonly transmitted on the Web. It uses a small, fixed set of markup tags to describe document portions. The small number of fixed tags simplifies document construction and makes it much easier to build applications. However, since the tags are fixed, HTML is not extensible and has very limited structure and validation capabilities. As electronic Web documents have become larger and more complex, it has become increasingly clear that HTML does not have the capabilities needed for large-scale commercial publishing.

In order to address the requirements of such large-scale commercial publishing and to enable the newly emerging technology of distributed document processing, an industry group called the World Wide Web Consortium has developed another markup language called the Extensible Markup Language (XML) for applications that require capabilities beyond those provided by HTML. Like HTML, XML is a simplified subset of SGML specially designed for Web applications and is easier to learn, use, and implement than full SGML. Unlike HTML, XML retains SGML advantages of extensibility, structure, and validation, but XML restricts the use of SGML constructs to ensure that defaults are available when access to certain components of the document is not currently possible over the Internet. XML also defines how Internet Uniform Resource Locators can be used to identify component parts of XML documents.

An XML document is composed of a series of entities or objects. Each entity can contain one or more logical elements and each element can have certain attributes or properties that describe the way in which it is to be processed. XML provides a formal syntax for describing the relationships between the entities, elements and attributes that make up an XML document. This syntax tells the computer how to recognize the component parts of each document.

XML uses paired markup tags to identify document components. In particular, the start and end of each logical element is clearly identified by entry of a start-tag before the element and an end-tag after the element. For example, the tags `<to>` and `</to>` could be used to identify the "recipient" element of a document in the following manner:

document text ... `<to>Recipient</to>` ... document text.

The form and composition of markup tags can be defined by users, but are often defined by a trade association or similar body in order to provide interoperability between users. In order to operate with a predefined set of tags, users need to know how the markup tags are delimited from normal text and the relationship between the

various elements. For example, in XML systems, elements and their attributes are entered between matched pairs of angle brackets (<...>), while entity references start with an ampersand and end with a semicolon (&...;). Because XML tag sets are based on the logical structure of the document, they are easy to read and understand.

5 Since different documents have different parts or components, it is not practical to predefine tags for all elements of all documents. Instead, documents can be classified into "types" which have certain elements. A document type definition (DTD) indicates which elements to expect in a document type and indicates whether each element found in the document is not allowed, allowed and required or allowed, but not
10 required. By defining the role of each document element in a DTD, it is possible to check that each element occurs in a valid place within the document. For example, an XML DTD allows a check to be made that a third-level heading is not entered without the existence of a second-level heading. Such a hierarchical check cannot be made with HTML. The DTD for a document is typically inserted into the document header and
15 each element is marked with an identifier such as <!ELEMENT>.

However, unlike SGML, XML does not require the presence of a DTD. If no DTD is available for a document, either because all or part of the DTD is not accessible over the Internet or because the document author failed to create the DTD, an XML system can assign a default definition for undeclared elements in the document.

20 XML provides a coding scheme that is flexible enough to describe nearly any logical text structure, such as letters, reports, memos, databases or dictionaries. However, XML does not specify how an XML-compliant data structure is to be stored and displayed, much less efficiently stored and displayed. Consequently, there is a need for a storage mechanism that can efficiently manipulate and store XML-compliant
25 documents.

SUMMARY OF THE INVENTION

In accordance with one embodiment of the invention, an in-memory storage manager represents XML-compliant documents as a collection of objects in memory.
30 The collection of objects allows the storage manager to manipulate the document, or

parts of the document with a consistent interface and to provide for features that are not available in conventional XML documents, such as element attributes with types other than text and documents that contain binary, rather than text, information. In addition, in the storage manager, the XML-compliant document is associated with a schema document (which is also an XML document) that defines the arrangement of the document elements and attributes. The storage manager can operate with conventional storage services to persist the XML-compliant document. Storage containers contain pieces of the document that can be quickly located by the storage manager.

In accordance with another embodiment, the storage manager also has predefined methods that allow it to access and manipulate elements and attributes of the document content in a consistent manner. For example, the schema data can be accessed and manipulated with the same methods used to access and manipulate the document content.

In accordance with yet another embodiment, the schema data associated with a document can contain a mapping between document elements and program code to be associated with each element. The storage manager further has methods for retrieving the code from the element tag. The retrieved code can then be invoked using attributes and content from the associated element and the element then acts like a conventional object.

In all embodiments, the storage manager provides dynamic, real-time data access to clients by multiple processes in multiple contexts. Synchronization among multiple processes accessing the same document is coordinated with event-driven queues and locks. The objects that are used to represent the document are constructed from common code found locally in each process. In addition, the data in the objects is also stored in memory local to each process. The local memories are synchronized by means of a distributed memory system that continually equates the data copies of the same element in different processes.

In still another embodiment, client-specified collections are managed by a separate collection manager. The collection manager maintains a data structure called a "waffle" that represents the XML data structures in tabular form. A record set engine

that is driven by user commands propagates a set of updates for a collection to the collection manager. Based on those updates, the collection manager updates index structures and may notify waffle users via the notification system. The waffle user may also navigate within the collection using cursors.

5

BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of the invention may be better understood by referring to the following description in conjunction with the accompanying drawings in which:

10 Figure 1 is a schematic diagram of a computer system on which the inventive storage manager system can run.

Figure 2 is a block schematic diagram illustrating the relationship of the in-memory storage manager and persistent storage.

45 Figure 3 is a block schematic diagram illustrating the representation of an XML document on the storage manager memory as a collection of objects.

Figure 4A is a block schematic diagram illustrating the components involved in binding code to XML elements.

Figure 4B is a flowchart showing the steps involved in retrieving program code bound to an element.

20 Figure 5 illustrates the relationship of XML text documents and binary sub-documents.

Figure 6 is a block schematic diagram illustrating the major internal parts of the storage manager in different processes.

Figure 7 illustrates the mechanism for synchronizing objects across processes.

25 Figure 8 is an illustration that shows the major control paths from the storage manager APIs through the major internal parts of the storage manager.

Figure 9 is an illustration of the storage manager interface constructed in accordance with an object-oriented implementation of the invention.

Figure 10 is an illustration of the interfaces constructed in accordance with an object-oriented implementation of the invention, that are defined by the storage manager and may be called during the processing of links or element RPCs.

Figure 11 is an illustration of the database and transaction interfaces constructed in accordance with an object-oriented implementation of the invention.

Figure 12 is an illustration of the document and element interfaces constructed in accordance with an object-oriented implementation of the invention.

Figure 13 is an illustration of the element communication and synchronization interfaces constructed in accordance with an object-oriented implementation of the invention.

Figure 14 is an illustration that shows the major control paths from the collection manager APIs through the major internal parts of the collection and storage managers.

Figure 15 is an illustration of the collection manager interfaces constructed in accordance with an object-oriented implementation of the invention.

DETAILED DESCRIPTION

Figure 1 illustrates the system architecture for an exemplary client computer 100, such as an IBM THINKPAD 600®, on which the disclosed document management system can be implemented. The exemplary computer system of Figure 1 is discussed only for descriptive purposes, however, and should not be considered a limitation of the invention. Although the description below may refer to terms commonly used in describing particular computer systems, the described concepts apply equally to other computer systems, including systems having architectures that are dissimilar to that shown in Figure 1 and also to devices with computers in them, such as game consoles or cable TV set-top boxes, which may not traditionally be thought of as computers.

The client computer 100 includes a central processing unit (CPU) 105, which may include a conventional microprocessor, random access memory (RAM) 110 for temporary storage of information, and read only memory (ROM) 115 for permanent storage of information. A memory controller 120 is provided for controlling system RAM

110. A bus controller 125 is provided for controlling bus 130, and an interrupt controller 135 is used for receiving and processing various interrupt signals from the other system components.

Mass storage may be provided by diskette 142, CD-ROM 147, or hard disk 152.

5 Data and software may be exchanged with client computer 100 via removable media, such as diskette 142 and CD-ROM 147. Diskette 142 is insertable into diskette drive 141, which is connected to bus 130 by controller 140. Similarly, CD-ROM 147 can be inserted into CD-ROM drive 146, which is connected to bus 130 by controller 145. Finally, the hard disk 152 is part of a fixed disk drive 151, which is connected to bus 130
10 by controller 150.

User input to the client computer 100 may be provided by a number of devices. For example, a keyboard 156 and a mouse 157 may be connected to bus 130 by keyboard and mouse controller 155. An audio transducer 196, which may act as both a microphone and a speaker, is connected to bus 130 by audio controller 197. It should
15 be obvious to those reasonably skilled in the art that other input devices, such as a pen and/or tablet and a microphone for voice input, may be connected to client computer 100 through bus 130 and an appropriate controller. DMA controller 160 is provided for performing direct memory access to system RAM 110. A visual display is generated by a video controller 165, which controls video display 170.

20 Client computer 100 also includes a network adapter 190 that allows the client computer 100 to be interconnected to a network 195 via a bus 191. The network 195, which may be a local area network (LAN), a wide area network (WAN), or the Internet, may utilize general-purpose communication lines that interconnect multiple network devices.

25 Client computer system 100 generally is controlled and coordinated by operating system software, such as the WINDOWS NT[®] operating system (available from Microsoft Corp., Redmond, WA). Among other computer system control functions, the operating system controls allocation of system resources and performs tasks such as process scheduling, memory management, networking and I/O services.

As illustrated in more detail in Figure 2, the storage manager 206 resides in RAM 200 (equivalent to RAM 110 in Figure 1) and provides an interface between an application program 202 which uses XML documents 228 and 230 and the persistent storage 208 in which the documents 228 and 230 are stored. The application 202 can interact with storage manager 206 by means of a consistent application programming interface 204 irregardless of the type of persistent storage 208 used to store the objects. Internally, the storage manager 206 represents each document 210, 218, as a hierarchical series of objects 212-216 and 220-224, respectively. The storage manager 206 can store the documents 210 and 218 in persistent storage 208 as schematically illustrated by arrow 226 using a variety of file systems, such as directory-based file services, object stores and relational file systems.

The inventive system operates with conventional XML files. A complete XML file normally consists of three components that are defined by specific markup tags. The first two components are optional, the last component is required, and the components are defined as follows:

1. An XML processing statement which identifies the version of XML being used, the way in which it is encoded, and whether it references other files or not. Such a statement takes the form:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
```

2. A document type definition (DTD) that defines the elements present in the file and their relationship. The DTD either contains formal markup tag declarations describing the type and content of the markup tags in the file in an internal subset (between square brackets) or references a file containing the relevant markup declarations (an external subset). This declaration has the form:

```
<!DOCTYPE Appl SYSTEM "app.dat">
```


3. A tagged document instance which consists of a root element, whose element type name must match the document type name in the document type declaration. All other markup elements are nested in the root element.

5 If all three components are present, and the document instance conforms to the document model defined in the DTD, the document is said to be "valid." If only the last component is present, and no formal document model is present, but each element is properly nested within its parent elements, and each attribute is specified as an attribute name followed by a value indicator (=) and a quoted string, document instance is said to be "well-formed." The inventive system can work with and generate well-formed XML documents.

10 Within the storage manager 206, XML documents are represented by means of data storage partitions which are collectively referred to by the name "Groove Document" to distinguish the representation from conventional XML documents. Each Groove document can be described by a DTD that formally identifies the relationships between the various elements that form the document. These DTDs follow the standard XML format. In addition, each Groove document has a definition, or schema, that describes the pattern of elements and attributes in the body of the document. XML version 1.0 does not support schemas. Therefore, in order to associate a Groove schema document with an XML data document, a special XML processing instruction containing a URI reference to the schema is inserted in the data document. This processing instruction has the form:

20 <?schema URI="groovedocument:///GrooveXSS/\$PersistRoot/sample.xml"?>

25 Some elements do not have, or require, content and act as placeholders that indicate where a certain process is to take place. A special form of tag is used in XML to indicate empty elements that do not have any contents, and therefore, have no end-tag. For example, a <ThumbnailBox> element is typically an empty element that acts

as a placeholder for an image embedded in a line of text and would have the following declaration within a DTD:

```
<!ELEMENT ThumbnailBox EMPTY >
```

Where elements can have variable forms, or need to be linked together, they can be given suitable attributes to specify the properties to be applied to them. These attributes are specified in a list. For example, it might be decided that the <ThumbnailBox> element could include a Location and Size attributes. A suitable attribute list declaration for such an attribute would be as follows:

```
<!ATTLIST ThumbnailBox
    Location    ENTITY    #REQUIRED
    Size        CDATA     #IMPLIED
>
```

This tells the computer that the <ThumbnailBox> element includes a required Location entity and may include a Size attribute. The keyword #IMPLIED indicates that it is permissible to omit the attribute in some instances of the <ThumbnailBox> element.

XML also permits custom definition statements similar to the #DEFINE statements used with some compilers. Commonly used definitions can be declared within the DTD as "entities." A typical entity definition could take the form:

```
<!ENTITY BinDoc3487 SYSTEM "../3487.gif" NDATA>
```

which defines a file location for the binary document "BinDoc3487." Once such a declaration has been made in the DTD, users can use a reference in place of the full value. For example, the <ThumbnailBox> element described previously could be specified as <ThumbnailBox Location=BinDoc3487 Size="Autosize"/>. An advantage of using this technique is that, should the defined value change at a later time, only the entity declaration in the DTD will need to be updated as the entity reference will automatically use the contents of the current declaration.

Within the storage manager, each document part is identified by a Uniform Resource Identifier (URI) which conforms to a standard format such as specified in RFC 2396. URIs can be absolute or relative, but relative URIs must be used only within the context of a base, absolute URI. When the document is stored in persistent storage, its parts may be identified by a different STORAGEURI that is assigned and managed by the particular file system in use.

In accordance with the principles of the invention, within each document part, in the storage manager internal memory is represented by a collection of objects. For example, separate elements in the XML document are represented as element objects in the storage manager. This results in a structure that is illustrated in Figure 3. In Figure 3, an illustrative XML document 300 is represented as a collection of objects in storage manager 302. In particular, the XML document 300 contains the conventional XML processing statement 304 which identifies the XML version, encoding and file references as discussed above. Document 300 also contains an XML processing statement 306 which identifies a schema document 320 in storage manager 302 which is associated with the document 300. The illustrative XML document also contains a set of hierarchical elements, including ElementA 308 which contains some text 318, ElementA contains ElementB 310 which has no text associated with it. ElementB also contains ElementC 312, which, in turn, contains two elements. Specifically, ElementC contains ElementD 314 that has an attribute (ID, with a value "foo") and ElementE 316.

In the storage manager 302, the elements, ElementA - ElementE, are represented as element objects arranged in a hierarchy. In particular, ElementA is represented by ElementA object 322. Each element object contains the text and attributes included in the corresponding XML element. Therefore, element object 322 contains the text 318. Similarly, ElementB 310 is represented by element object 324 and elements ElementC, ElementD and ElementE are represented by objects 326, 328 and 330, respectively. Element object 328, which represents element ElementD, also includes the attribute ID that is included in the corresponding element. Each element object references its child element objects by means of database pointers (indicated by arrows between the objects) into order to arrange the element objects into a hierarchy.

There may also be attribute indices, such as index 332 that indexes the ID attribute in element object 328.

The representation of the XML document 300 by means of an object collection allows the storage manager 302 to manipulate its internal representation of the document 300 with a consistent interface that is discussed in detail below. The storage manager 302 can also provide features that are not available in conventional XML documents, such as collection services that are available via a collection manager that is also discussed in detail below.

As described above, Groove documents that contain XML data may have a definition, or schema document, that describes the pattern of elements and attributes in the body of the document. The schema document is stored in a distinct XML document identified by a URI. The schema document has a standard XML DTD definition, called the meta-schema, which is shown below:

```
<!-- The Document element is the root element in the schema -->
<!ELEMENT Document (Registry*, AttrGroup*, ElementDecl*)>
<!ATTLIST Document
    URL          CDATA          #REQUIRED
>

<!ELEMENT Registry TagToProgID*>

<!ELEMENT TagToProgID EMPTY>
<!ATTLIST TagToProgID
    Tag          CDATA          #REQUIRED
    ProgID       CDATA          #REQUIRED
>

<!ELEMENT AttrGroup AttrDef*>

<!ELEMENT AttrDef EMPTY>
<!ATTLIST AttrDef
    Name          CDATA          #REQUIRED
    Type          CDATA          #REQUIRED
    Index         CDATA          #IMPLIED
    DefaultValue  CDATA          #IMPLIED
>
```

```

<!ELEMENT ElementDecl (ElementDecl* | AttrGroup | ElementRef*)>
<!ATTLIST ElementDecl
    Name      CDATA      #REQUIRED

```

```

>

<!ELEMENT ElementRef EMPTY>
<!ATTLIST ElementRef
    Ref      CDATA      #REQUIRED

```

Each of the elements in the schema defines information used by the storage manager while processing the document. The “Registry” section forms an XML representation of a two-column table that maps XML element tags to Windows ProgIDs. (In the Common Object Model (COM) developed by Microsoft Corporation, a ProgID is a text name for an object that, in the COM system, is “bound” to, or associated with, a section of program code. The mapping between a given ProgID and the program code, which is stored in a library, is specified in a definition area such as the Windows™ registry.)

This arrangement is shown in Figure 4A that illustrates an XML document 400 and its related schema document 402. Both of these documents are resident in the storage manager 406 and would actually be represented by objects as shown in Figure 3. However, in Figure 4, the documents have been represented in conventional XML format for clarity. Figure 4 shows the storage manager operational in a Windows™ environment that uses objects constructed in accordance with the Common Object Model (COM) developed by the Microsoft Corporation, Redmond, Washington, however, the same principles apply in other operating system environments.

XML document 400 includes the normal XML processing statement 414 that identifies the XML version, encoding and file references. A schema XML processing statement 416 references the schema document 402 which schema document is associated with document 400 and has the name “urn:groove.net:sample.xml” defined by name statement 426. It also includes a root element 418 which defines a name “doc.xml” and the “g” XML namespace which is defined as “urn:groove.net”

Document 400 has three other elements, including element 420 defined by tag "urn:groove.net:AAA", element 422 defined by tag "urn:groove.net:BBB" and element 424 defined by tag "urn:groove.net:NoCode". Element 424 is a simple element that has no corresponding bound code and no corresponding tag-to-ProgID mapping in the schema document 402.

Within the "registry" section defined by tag 428, the schema document 402 has two element-to-COM ProgID mappings defined. One mapping is defined for elements with the tag "urn:groove.net:AAA" and one for elements with the tag "urn:groove.net:BBB." The bound code is accessed when the client application 404 invokes a method "OpenBoundCode()." The syntax for this invocation is given in Table 15 below and the steps involved are illustrated in Figure 4B. Invoking the OpenBoundCode() method on a simple element, such as element 424 generates an exception. The process of retrieving the bound code starts in step 434 and proceeds to step 436 in which the OpenBoundCode() is invoked. Invoking the OpenBoundCode() method on an element with the element tag "urn:groove.net:AAA" causes the storage manager 406 to consult the registry element 428 in the schema document 602 with the element tag as set forth in step 438. From section 430, the storage manager retrieves the ProgID "Groove.Command" as indicated in step 440. In step 442, the storage manager calls the COM manager 408 in instructs it to create an object with this ProgID. In a conventional, well-known manner, in step 444, the COM manager translates the ProgID to a CSLID using a key in the Windows Registry 410. In step 446, the COM manager uses the CSLID to find a dynamically loadable library (DLL) file in the code database 412 that has the code for the object. Finally, in step 448, the COM manager creates the object and returns an interface pointer for the object to the storage manager 406 which, in turn, returns the pointer to the client application 404. The routine then finishes in step 450. The client application 404 can then use the pointer to invoke methods in the code that use attributes and content in the associated element. The element then behaves like any other COM object. A similar process occurs if the OpenBoundCode() method is invoked on elements with the tag "urn:groove.net:BBB."

The “AttrGroup” section defines non-XML characteristics for attributes. An attribute’s data type can be defined as some type other than text and the attribute may be indexed to facilitate fast retrieval of the elements that containing it.

The “ElementDecl” section provides a form of element definition similar to the DTD <!ELEMENT> declaration, but allows for extended attribute characteristics and the definition of non-containment element references.

The following example shows the sample portions of a schema document for an XML document that defines a “telespace” that is previously described.

```
<groove:Document URL="TelespaceSchema.xml"
  xmlns:groove="urn:groove.net:schema.1">
  <groove:Registry>
    <groove:TagToProgID groove:Tag="g:Command"
      groove:ProgID="Groove.Command"/>
    <groove:TagToProgID groove:Tag="groove:PropertySetChanged"
      groove:ProgID="Groove.PropSetChangeAdvise"/>
  </groove:Registry>
  <groove:AttrGroup>
    <groove:AttrDef Name="ID" Index="true"/>
    <!-- KEY EXCHANGE ATTRIBUTES -->
    <groove:AttrDef Name="NKey" Type="Binary"/>
    <groove:AttrDef Name="ReKeyld" Type="String"/>
    <groove:AttrDef Name="T" Type="String"/>
    <!-- AUTHENTICATION ATTRIBUTES -->
    <groove:AttrDef Name="MAC" Type="Binary"/>
    <groove:AttrDef Name="Sig" Type="Binary"/>
    <!-- ENCRYPTION ATTRIBUTES -->
    <groove:AttrDef Name="IV" Type="Binary"/>
    <groove:AttrDef Name="EC" Type="Binary"/>
    <!-- XML Wrapper Attributes -->
    <groove:AttrDef Name="Rows" Type="Long"/>
    <groove:AttrDef Name="Cols" Type="Long"/>
    <groove:AttrDef Name="Items" Type="Long"/>
    <groove:AttrDef Name="ItemID" Type="Bool" Index="true"/>
  </groove:AttrGroup>
  <groove:ElementDecl Name="groove:Telespace">
    <AttrGroup>
      <AttrDef Name="Persist" DefaultValue="True" Type="Bool"/>
      <AttrDef Name="Access" DefaultValue="Identity" Type="String"/>
    </AttrGroup>
```

```

        <ElementRef Element="Dynamics"/>
        <ElementRef Element="Members"/>
    </groove:ElementDecl>
</groove:Document>

```

5

In this example, there are two entries in the Tag to ProgID mapping table. The first maps the tag "g:Command" (which, using XML namespace expansion, is "urn:groove.net.schema.1:Command") to the ProgID "Groove.Command." In the section defining attributes, the "ID" attribute is indexed, the data type of the NKey attribute is binary, and so on.

10

This schema data is represented by element objects and can be accessed and manipulated by the same storage manager element and attribute interface methods used to manipulate documents as described in detail below. In particular, the information that describes a document can be manipulated using the same interfaces that are used for manipulating the document content.

15

In accordance with another aspect of the invention, sub-documents can be associated with a primary document. Any document may be a sub-document of a given document. If a document contains a sub-document reference to another document, then the referenced document is a sub-document. If two documents contain sub-document references to each other, then each document is a sub-document of the other document. Each sub-document is referenced from the primary document with conventional XML XLink language, which is described in detail at <http://www.w3.org/TR/xlink/>. Links may also establish a relationship between an all-text XML document and a binary sub-document. Binary documents do not have links to any kind of sub-document. If the link is to a document fragment, a subdocument relationship is established with the document that contains the fragment. The relationship of documents and sub-documents is illustrated in Figure 5.

20

25

30

For example, main document 500 contains links 502 which include a link, represented by arrow 510, to document 504 and a link, represented by arrow 508, to a binary document 506. Documents 504 and 506 are thus sub-documents of document 500. Document 504, in turn, contains links 512 which include a link, represented by

arrow 514 to document 516 with content 518. Document 516 is a sub-document of document 500. Document 506 contains binary content 520 and, therefore, cannot have links to sub-documents.

Sub-document links follow the standard definition for simple links. An exemplary element definition of a link is as follows:

```
<!ELEMENT GrooveLink ANY>
<!ATTLIST GrooveLink
    xml:link          CDATA          #FIXED "simple"
    href              CDATA          #REQUIRED
    role              CDATA          #IMPLIED "sub-document"
    title             CDATA          #IMPLIED
    show              (parsed|replace|new) #IMPLIED
    actuate            (auto|user)      #IMPLIED
    serialize          (byvalue|byreference|ignored) #IMPLIED
    behavior           CDATA          #IMPLIED
    content-role       CDATA          #IMPLIED
    content-title      CDATA          #IMPLIED
    inline             (true|false)    #IMPLIED "true"
>
```

It is also possible to establish a sub-document relationship without using the above definition by adding to a document an XML link which has an xml:link attribute with a value "simple", and a href attribute. Such a link will establish a sub-document relationship to the document identified by a URI value in the href attribute.

Given the relationships from a document to its sub-documents, it is possible to make a copy of an arbitrary set of documents and sub-documents. Within a single storage service, it may be possible to directly perform such a copy. To cross storage services or to send multiple documents to another machine, the entire hierarchy of such documents must be "describable" in a serialized fashion. The inventive Storage Manager serializes multiple documents to a text representation conforming to the specification of MIME Encapsulation of Aggregate documents, such as HTML (MHTML) which is described in detail at <ftp://ftp.isi.edu/in-notes/rfc2557.txt/>.

The following data stream fragment is an example of a document and a referenced sub-document as they would appear in an MHTML character stream. In the

example, "SP" means one space is present and "CRLF" represents a carriage return-line feed ASCII character pair. All other characters are transmitted literally. The MIME version header has the normal MIME version and the Groove protocol version is in a RFC822 comment. The comment is just the word "Groove" followed by an integer. The boundary separator string is unique, so a system that parsed the MIME, and then each body part, will work correctly. The serialized XML text is illustrated in UTF-8 format, but it could also be transmitted in WBXML format. The XML document has a XML prefix, which includes the version and character encoding. The binary document is encoded in base64.

```

10  MIME-Version: SP 1.0 SP (Groove SP 2) CRLF
    Content-Type: SP multipart/related; SP boundary="<[&&&]>" CRLF
    CRLF
    --<[&&&]>>Content-Type: SP text/XML; SP charset="UTF-8"
    <?xml version="1.0" encoding='utf-8'?>
15  <rootelement>
    .
    .
    .
    </rootelement> CRLF
20  CRLF
    --<[&&&]>>
    Content-ID: SP <URI> CRLF
    Content-Type: SP application/octet-stream CRLF
    Content-Transfer-Encoding: base64 CRLF
25  CRLF
    R0lGODIhdQAgAPcAAP////////zP//mf//Zv//M//AP/M//MzP/Mmf/MZv/MM//MAP+Z//+Z
    zP+Zmf+ZZv+ZM//+ZAP9m//9mzP9mmf9mZv9mM/9mAP8z//8zzP8zmf8zZv8zM/8zAP8A//8A
    zP8Amf8AZv8AM/8AAMz//8z/zMz/mcz/Zsz/M8z/AMzM/8zMzMzMmczMZszMM8zMAMyZ/8yZ
    zMyZmcyZZsyZM8yZAMxm/8xmzMxmcmxmZsxmM8xmAMwz/8wzzMwzmcwzZswzM8wzAMwA/8wA
30  zMwAmcwAZswAM8wAAJn//5n/zJn/mZn/Zpn/M5n/AJnM/5nMzJnMmZnMZpnMM5nMAJmZ/5mZ
    OG/qTMnzJUWQHOMKHUq0KEagRpMqXaoUaU6dG2IKIOqRKtOkTq9q3VrV5sd/XMOKZZp1rNmz
    GsuiXct2hNq2cMVmXdkzZ12LLe/ehYrXpsy/MPUGHvw04IzCdhFbzayMd+aUxsnnrzTq1uw
    cTN3tVrxrebPWDGDHr3UM+nTHE2jXn1RNeVXEI3Dfi179urDJrte5BzVcknNhyNHZiyzJnGv
    uWMuppu7uHLkyV1Kxe1ccOGZ0Cn/xshcu8/K2g2LQ8bJGPJj4eh3+/WNHb118PAAtBn8aXTrn
35  6s7tl2QP9b399fhNN55tbe31FYEITIRbgqAtyCBwAz5l20MUVmjhhRgyFBAAOw==
    --<[&&&]>>--

```

Unlike most XML processors, such as document editors or Internet browsers, the storage manager provides for concurrent document operations. Documents may be concurrently searched, elements may be concurrently created, deleted, updated, or

moved. Copies of element hierarchies may be moved from one document to another. In most XML processors, all of the updates to a document are driven by a single user, who is usually controlling a single thread within a single process on a single computer.

The storage manager maintains XML document integrity among many users

5 updating the same document, using multiple threads in multiple processes. In a preferred embodiment, all of the updates occur on a single computer, but, using other different, conventional inter-processor communication mechanisms, other operational embodiments are possible. Figure 6 shows the basic structure of the storage manager and illustrates how it isolates application programs from cross-process communication
10 issues. For example, two separate processes 600 and 602 may be operating concurrently in the same computer or in different computers. Process 600 is a "home" process as described below, while process 602 is another process designated as Process N. Within process 600, a multi-threaded client application program 606 is operating and within process 602, a multi-threaded client application program 616 is
15 operating.

Each application program 606 and 616 interfaces with a storage manager designated as 605 and 615, respectively. In process 600, the storage manager comprises a storage manager interface layer 608 which is used by application program 608 to control and interface with the storage manager. It comprises the database,
20 document, element and schema objects that are actually manipulated by the application. The API exported by this layer is discussed in detail below. The storage manager 605 also includes distributed virtual object (DVO) database methods 610, DVO methods for fundamental data types 612, DVO common system methods 609 and distributed shared memory 614. Similarly, the storage manager operating in process
25 602 includes transaction layer 618, DVO database methods 620, DVO methods for fundamental data types 622, DVO common system methods 617 and distributed shared memory 624.

The two processes 600 and 602 communicate via a conventional message passing protocol or inter-process communication (IPC) system 604. For processes that
30 run in a single computer, such a system can be implemented in the Windows® operating

system by means of shared memory buffers. If the processes are running in separate computers, another message passing protocol, such as TCP/IP, can be used. Other conventional messaging or communications systems can also be used without modifying the operation of the invention. However, as is shown in Figure 6, application programs 606 and 616 do not directly interact with the message passing system 604. Instead, the application programs 606 and 616 interact with storage managers 605 and 615, respectively, and storage managers 605 and 615 interact with the message passing system 604 via a distributed shared memory (DSM) system of which DSM systems 614 and 624 are a part.

A number of well-known DSM systems exist and are suitable for use with the invention. In accordance with a preferred embodiment, the DSM system used with the storage manager is called a C Region Library (CRL) system. The CRL system is an all-software distributed shared memory system intended for use on message-passing multi-computers and distributed systems. A CRL system and code for implementing such as system is described in detail in an article entitled "CRL: High-Performance All-Software Distributed Memory System", K. L. Johnson, M.F. Kaashoek and D. A. Wallach, *Proceedings of the Fifteenth Symposium on Operating Systems Principles*, ACM, December 1995; and "CRL version 1.0 User Documentation", K.L. Johnson, J. Adler and S. K. Gupta, , MIT Laboratory for Computer Science, Cambridge, MA 02139, August 1995. Both articles are available at Web address <http://www.pdos.lcs.mit.edu/crl/>.

Parallel applications built on top of the CRL, such as the storage manager, share data through memory "regions." Each region is an arbitrarily sized, contiguous area of memory . Regions of shared memory are created, mapped in other processes, unmapped, and destroyed by various functions of the DSM system. The DSM system used in the present invention provides a super-set of the functions that are used in the CRL DSM system. Users of memory regions synchronize their access by declaring to the DSM when they need to read from, or write to, a region, and then, after using a region, declaring the read or write complete. The effects of write operations are not propagated to other processes sharing the region until those processes declare their

need for it. In addition to the basic shared memory and synchronization operations, DSM provides error handling and reliability with transactions. The full interface to inventive DSM is shown in Table 1.

5

TABLE 1

DSM Method	Description
AddNotification(DSMRgn* i_pRgn, const IgrooveManualResetEvent * i_pEvent);	Adds a local event that will be signaled with the data in the region changes.
Close();	Shuts down the DSM. There must be no mapped regions at this client.
Create(UINT4 i_Size, INT4 i_CallbackParam, INCAAddress i_InitialOwner, DSMRId & io_RId, DSMRgn * & o_pRgn, void * & o_pData);	Creates a new region. It also atomically maps the new region and initiates a StartWrite on the new region if Size is non-zero. Size is the initial size of the data in the new region. RId is identifier of the new region. pRgn is the new region if Size is non-zero.
AddDatabase(UINT2 i_DatabaseNumber);	Adds a new database to the region mapping tables.
DatabaseFlushNotify(UINT2 i_DatabaseNumber, TimeMillis i_StartTime);	Cleans up unused region resources.
Destroy(DSMRId& i_RId);	Destroys an existing region entirely. RId is a valid identifier of the region to be destroyed.
EndRead(DSMRgn* i_pRgn);	Closes a read operation on the region's data. pRgn is the valid region.
EndWrite(DSMRgn* i_pRgn);	Closes a write operation on the region's

Flush(DSMRgn* i_pRgn);

GetSize(DSMRgn* i_pRgn);

Init(CBSTR i_BroadcastGroup,
DSMRgnMapCallback * i_pCallback =
NULL, void * i_pCallbackParam = NULL,
BOOL * o_pMasterClient = NULL, UINT4
i_WaitTimeOut = 1000, UINT4 i_URCSize
= 1<<10, INCAAddress * o_pAddress =
NULL);

Map(const DSMRId& i_RId, INT4
i_CallbackParam, BOOL i_InitialOwner);

RemoveDatabase(UINT2
i_DatabaseNumber);

RemoveNotification(DSMRgn* i_pRgn,
const IGrooveManualResetEvent *
i_pEvent);

Resize(DSMRgn* i_pRgn, UINT4 i_Size);

data. pRgn is the valid region.

Flushes the region from this client's local cache to the region's home client. pRgn is the valid region.

Returns the size(number of bytes) of the given valid region. pRgn is the valid region.

Initializes the DSM. BroadcastGroup is the name of the group in which this DSM client belongs. URCSize is the size of the Unmapped Regions Cache. PAddress is the Inter-node Communication Address of this DSM client. pMasterClient specifies whether this DSM client is the Master(First) client.

Maps the region to this client's memory space. RId is a valid identifier of the region to be mapped.

Removes the specified database from the region mapping tables.

Removes interest in changes to data in a region.

Resizes the given valid region while maintaining the original data(which may be truncated if the size is decreased). pRgn is the valid region. Size is the new size.

GetRId(const DSMRgn* i_pRgn);

Returns the identifier for the given valid region. pRgn is the valid region.

SignalNotification(DSMRgn* i_pRgn);

Sets the signal that notification has occurred.

StartRead(DSMRgn* i_pRgn, INT4
i_CallbackParam, void * & o_pData);

Initiates a read operation on the region's data. RgnStartRead (or RgnStartWrite) must be called before the data can be read. pRgn is the valid region.

StartTransactionRead(DSMRgn* i_pRgn,
INT4 i_CallbackParam, void * & o_pData);

Initiates a transactional read operation on the region's data. RgnStartRead (or RgnStartWrite) must be called before the data can be read. pRgn is the valid region.

StartTransactionWrite(DSMRgn* i_pRgn,
INT4 i_CallbackParam, void * & o_pData);

Initiates a transactional write operation on the region's data. RgnStartWrite must be called before the data can be modified. pRgn is the valid region.

StartWrite(DSMRgn* i_pRgn, INT4
i_CallbackParam, void * & o_pData);

Initiates a write operation on the region's data. RgnStartWrite must be called before the data can be modified. pRgn is the valid region.

Unmap(DSMRgn* & io_pRgn);

Unmaps the region from this client's memory space. pRgn is the valid region to be unmapped.

Each storage manager 605 and 615 comprises a DSM node that uses one or more DSM regions (not shown in Figure 6) located in the address space of the corresponding process 600, 602. These regions contain DVO objects and classes that can be used to represent documents, elements and schema of the XML data that is

managed by the storage manager. Portions of documents, usually elements and index sections, are wholly contained within a region. Although the DSM system provides a conceptually uniform node space for sharing regions, there are issues that result in the need to single out a specific node or process to perform special tasks.

5 Consequently, within the DSM synchronization protocol, a single node is identified as a “home node” for each region. Within the many processes running the storage manager on a single computer, one process, called the “home process”, is the process that performs all disk I/O operations. To reduce the amount of data movement between processes, the home process is the home node for all regions. Other
10 implementations are possible, in which any node may be the home for any region and any process may perform disk I/O. However, for personal computers with a single disk drive, allowing multiple processes to perform disk I/O introduces the need for I/O synchronization while not alleviating the main performance bottleneck, which is the single disk.

15 In accordance with the DSM operation, if a process has the most recent copy of a region, then it can read and write into the region. Otherwise, the process must request the most-recent copy from the home process before it can read and write in the region. Each DSM system 614, 624 interfaces with the message passing system 604 via an interface layer called an internode communication layer (615, 625) which isolates
20 the DVM system from the underlying transport mechanism. It contains methods that send messages to a broadcast group, and manipulate addresses for the corresponding process and the home process.

25 The inventive storage manager uses shared objects as the basis for XML objects. Many systems exist for sharing objects across processes and computers. One such object-sharing model is based on the use of the shared memory facilities provided by an operating system. One of the biggest drawbacks of such a shared memory model is unreliability due to memory write failures that impact the integrity of other processes. For example, if one process is in the process of updating the state of an object and the process fails before setting the object to a known good state, other processes will either
30 see the object in an invalid state or may be blocked indefinitely waiting for the failed

process to release its synchronization locks. The shared memory model also suffers from the locality constraints of shared memory in a tightly coupled multi-computer – it provides no way to share objects over a network.

Another model that provides distributed object sharing and remote method invocation is the basis for the distributed object management facilities in Java or the Object Management Group's CORBA system. Although providing the ability to share objects over a computer network, clients of such systems need to be aware of whether an object is local or remote – objects are not location independent. Performance is another drawback of this approach. All operations on an object need to be transmitted to the object server, since the server contains the only copy of the object state and serves as the synchronization point for that data.

In order to overcome these drawbacks, the inventive storage manager uses a distributed virtual object (DVO) system to provide the primitive data types that XML object types are built upon. The DVO system also provides its callers with the illusion that all data is reliably contained in one process on a single computer node, even though the data may be in multiple processes on many computers or may truly be just in one process on a single computer node.

The DVO object-sharing model is shown in Figure 7. All processes, on all computers, that are sharing an object have the same method code. For example, process 700 and process 702 in Figure 7 have copies of the same object. Thus, each of processes 700 and 702 has a copy of the same method code 704 and 706 in the respective process address space. The volatile data state for an object is stored in DSM regions. Thus, the object data 708 for the object copy in process 700 is stored in region 710 in the address space of process 700. Similarly, the object data 712 for the object copy in process 702 is stored in region 714 in the address space of process 702. Object methods synchronize their access to the object's data by using the DSM synchronization functions that synchronize the regions as illustrated by arrow 716. In this manner, DVO objects are location independent, failures are contained within a single process, and multiple changes to a local object do not require data movement across the inter-node transport.

The DVO system provides basic objects that may be used as building blocks to manage XML documents for the storage manager and is divided into three functional pieces. The DVO database 610 contains objects that handle the DVO local context in each process and the shared tables that contain information about open databases and documents contained within those databases. In DVO, “databases” are conceptual storage containers and may channel objects that are ultimately stored in any kind of storage service 609. DVO documents are associated with XML or binary documents, which are visible to a client of the storage manager. DVO documents are also used to contain the indices and metadata associated with a collection.

DVO types 612 is a set of object classes that can be used within DVO documents to implement higher-level data model constructs. DVO types range from simple data containment objects through complex, scalable index structures. Each DVO type is implemented with two classes – one is a “non-shared class” that uses memory pointers in object references and the other is a “shared class” that uses logical addresses, called database pointers, for object references. The “shared class” has two sub-forms – one is the representation of the object in a shared DSM region and the other is the representation of the object stored on-disk in an object store database. The DVO system 607 provides methods to transfer objects between their shared and non-shared implementations.

The different DVO types are shown in Table 2.

TABLE 2

DVO Type	Description
Binary Document	A kind of document that handles binary data.
B-tree Index	The type of the root of a b-tree index. It contains a description of the index, as well as the address of the root index node.
Btree Node	A piece of a Btree index which can contain variable numbers of records, sorted by one or more keys.

Collection Document	A kind of document that handles Collection documents. In addition to the Document methods, it has methods to handle the collection descriptor, indices within the collection, and read marks.
Document	The base type from which the other document types inherit common methods, such as Open, Close, Create, and Write.
Extendible Hashing	A type implementation of extendible hashing, as defined in "Extendible Hashing - A Fast Access Method for Dynamic Files", Ronald Fagin, Jürg Nievergelt, Nicholas Pippenger, H. Raymond Strong. <i>ACM Transactions on Database Systems</i> 4(3), pages 315-344, 1979.
FlatCollectionDocument	A specific kind of CollectionDocument used in shared regions.
FlatDocument	A specific kind of XMLDocument used in shared regions.
FlatNode	A specific kind of Node used in shared regions.
Node	The type used to store XML elements. It has methods to manage the element name, the element's parent, element content, element attributes, links to other elements, and change notifications.
Ordered Bucket	A kind of index which supports key ordered sorting (integer, double, string)
Ordered Index	A type that provides a collated data vector. It has methods for adding, removing, and changing key/data pairs, managing index cursors, and managing parent and sub-indicies.
Ordered Index Types	Data types, called records and fields, that can be stored in ordered indices.

Ordinal Ordered Index	A kind of index that support ordinal addressing. It is conceptually similar to vector that allows any entry to be addressed by position (e.g., vec[14]). In addition to the index methods, it has methods to move entries to specific positions within the index.
Red-Black Index	A kind of ordered index that implements balancing using the red-black binary tree algorithm.
W32BinaryDocument	A specific kind of binary document for 32-bit Windows platforms.
XML Document	A kind of document that handles XML documents. In addition to the Document methods, it has methods to handle schemas and indexes.

The DVO system 607 objects isolate the upper levels of DVO from physical storage and process locality issues. The DVO system objects use DSM for invoking and handling requests to and from the home process. Requests include operations such as opening, closing, and deleting a database, finding documents in a database, and opening, closing, deleting, and writing database documents. The DVO system 607 in the master process 600 can also retrieve DVO objects from a storage service 609. A storage service, such as service 609, is a utility program that stores and retrieves information from a persistent medium and is responsible for the physical integrity of a container, database or file. It ensures that all updates are durable and that all internal data structures (e.g., redirection tables, space allocation maps) are always consistent on disk. Other processes, such as process 602 cannot access the storage service 609 directly, but can access the system indirectly via its DSM regions 624.

The storage manager 605 can operate with different types of physical storage systems, including container or object stores, stream file systems and ZIP files. In order to achieve atomic commits, the object store storage service can be implemented using page-oriented input/output operations and a ping-pong shadow page table.

Individual storage manager methods are atomic. Multiple storage manager operations, even operations on different documents, may be grouped into “transactions.” Transactions not only protect XML data integrity, but they also improve performance because they enable the storage manager to reduce the number of region lock operations and reduce the amount of data movement over the message passing system.

The storage manager supports both read-write and read-only transactions built on DSM synchronization primitives described in the DSM documentation referenced above, which primitives insure consistency in multiple processes or computers. Read-write transactions provide for the atomicity and consistency of a set of database read and write operations. Each region that is changed as part of a transaction will be kept in a “locked” state until the transaction is committed or aborted. This prevents operations that are not part of the transaction from seeing the changes. Further, each transaction stores a “before image” of the regions it modifies so that, if the transaction is aborted (as a result of an explicit API call or an exception), the effects of the transaction can be undone. Depending on the performance requirements, an alternative implementation would write undo information rather than storing the full “before image.” A read-only transaction uses the same interface as a read-write transaction. A read-only transaction ensures that multiple read operations are consistent. Like other transactions, it uses DSM functions to keep all read regions in a “read state” until it is finished.

In addition, checkpoints can be used to ensure that changes are persistent and provide durability for storage manager operations. A checkpoint may be performed at any time. Checkpoints are used in conjunction with data recovery logging. All operations write “redo” information to a sequential recovery log file when they are committed. When the checkpoint is committed, the recovery log file will be flushed to persistent storage and will ensure that the operations can be recovered. Since transactions do not write “redo” information until they are committed, if a checkpoint operation is commenced in the middle of a transaction, the transaction operations will not be flushed.

Transactions are scoped to a thread and a database. Once a transaction is started on a thread for a particular database, that transaction will be automatically used for all subsequent storage manager operations on that database and thread. An extension of conventional operating system threads is used, so that transactions

5 correctly handle calls that need to be marshaled to other threads, for example, a user interface thread, using the Groove system's simple marshaler. Storage manager calls made on a thread and database that doesn't have a transaction started will cause the storage manager to create a "default transaction" that will be committed just before the call ends. Alternatively, starting a new transaction on a thread and database that

10 already has an existing transaction in progress will cause the new transaction to automatically "nest" in the existing transaction. Nested transactions provide the ability to roll back the system within the outer transaction. In particular, inner, nested transactions are not finally committed until the outermost transaction is committed. For example, if a nested transaction is committed, but the containing transaction is later

15 aborted, the nested transaction will be aborted.

In a preferred embodiment of the invention, the storage manager is implemented in an object-oriented environment. Accordingly, both the storage manager itself and all of the document components, such as documents, elements, entities, etc. are implemented as objects. These objects, their interface, the underlying structure and the

20 API used to interface with the storage manager are illustrated in Figure 8. The API is described in more detail in connection with Figures 9-11. Referring to Figure 8, the storage manager provides shared access to documents, via the document manipulation API 802, but, in order to enable a full programming model for client applications, additional communication and synchronization operations are provided, within the

25 context of a document. For example, the storage manager provides queued element operations, which enable one process to send an element to another process via the Queue API 804. Elements can be sent by value (a copy of the whole element) or by reference to the element. Synchronization operations are also provided to allow one or more threads to wait for an element to be enqueued to a given queue. The storage

manager also provides RPC-style element communication and synchronization, via the RPC API 804.

Other client components may need to be aware of when documents are created in or deleted from storage manager. Accordingly, the storage manager provides an interface to an interest-based notification system for those client components via notification API 800. The notification system 806 provides notifications to client components that have registered an interest when a document is created or deleted.

Document data is represented by a collection of objects including database objects, document objects, element objects and schema objects 808. The objects can be directly manipulated by means of the document manipulation API 802.

The document related objects 808 are actually implemented by the distributed virtual object system 810 that was discussed in detail above. The distributed virtual object system 810 can also be manipulated by element queue and RPC objects 812 under control of the queue and RPC API 804.

The distributed virtual object system 810 communicates with the distributed shared memory via interface 814 and communicates with the logging operations via interface 816. Similarly, the distributed virtual object system can interact with the storage services via interface 818.

The following is a description of the interfaces for each of the objects used to implement a preferred embodiment of the inventive storage manager. These object are designed in accordance with the Common Object Model (COM) promulgated by Microsoft Corporation, Redmond, Washington, and can be manipulated in memory as COM objects. However, COM is just one object model and one set of interface methodologies. The invention could also be implemented using other styles of interface and object models, including but not limited to the Java and CORBA object models.

Figure 9 illustrates object interfaces for a storage manager object. An interface 900 (IGrooveStorageManager) encapsulates the basic framework for the storage manager. This interface is a subclass of an IDispatch interface which is a common

class defined by the COM model. Table 3 defines the methods included in the storage manager interface.

TABLE 3

Interface IGrooveStorageManager : IDispatch

CreateDatabase (BSTR i_DatabaseURI, VARIANT_BOOL i_Temporary, VARIANT_BOOL i_SingleProcess, IUnknown * i_pSecurityContext, VARIANT_BOOL i_CreateOnCheckpoint, IgrooveDatabase ** o_ppDatabase);	Creates a database. A database can be either temporary or permanent, and single or multi-process. The DatabaseURI specifies the location of the database.
CreateOrOpenDatabase (BSTR i_DatabaseURI, VARIANT_BOOL i_Temporary, VARIANT_BOOL i_SingleProcess, IUnknown * i_pSecurityContext, VARIANT_BOOL i_CreateOnCheckpoint, VARIANT_BOOL * o_pCreated, IgrooveDatabase ** o_ppDatabase);	Creates a new database or opens an existing database.
CreateTemporaryElement (BSTR i_Name, IUnknown * i_pParent, IgrooveElement ** o_ppElement);	Creates a temporary element.

CreateTemporaryXMLDocument (BSTR i_NamePrefix, BSTR i_SchemaURI, IUnknown* i_pAdditionalSchemaURIs, IgrooveXMLDocument ** o_ppXMLDocument);	Creates an empty temporary document with a unique URI
CreateTransform (BSTR i_CollectionDescriptorURI, BSTR i_SecondaryDescriptorURI, BSTR i_CollectionDescriptorName, IgrooveTransform ** o_ppTransform);	Creates a transformation interface.
DeleteDatabase (BSTR i_DatabaseURI);	Deletes a database.
IsHomeProcess (VARIANT_BOOL * o_pHomeProcess);	Determine whether we are the home process
OpenCrossProcessSemaphore (BSTR i_Name, VARIANT_BOOL i_Reentrant, IgrooveCrossProcessSemaphore ** o_ppSemaphore);	Creates a semaphore object that can be used to synchronize activity in different processes. If the semaphore is not Reentrant, repeated attempts to lock the semaphore within the same thread and process will block.
OpenDatabase (BSTR i_DatabaseURI, VARIANT_BOOL i_SingleProcess, IUnknown * i_pSecurityContext, IgrooveDatabase ** o_ppDatabase);	Open an existing database.
OpenDatabaseURIEnum(IGrooveBST REnum ** o_ppDatabaseURI);	Returns an Enumeration of the databases that are currently open.

Another interface 902 (IGrooveStorageURISyntax) is used by a client of a storage manager that needs to perform operations on parts of standard names, which are in the form of Uniform Resource Identifiers (URIs). Table 4 includes the methods for the IGrooveStorageURISyntax interface.

5

TABLE 4

Interface IGrooveStorageURISyntax : IDispatch

BuildDatabaseURI (BSTR i_ServiceName, BSTR i_DatabasePath, VARIANT_BOOL i_Relative, BSTR *o_pURI);	Builds a database URI from its pieces.
BuildDocumentURI (BSTR i_ServiceName, BSTR i_DatabasePath, BSTR i_DocumentName, VARIANT_BOOL i_Relative, BSTR * o_pURI);	Builds a document URI from its pieces.
MakeAbsolute (BSTR i_RelativeURI, BSTR * o_pAbsoluteURI);	Given a relative URI within the scope of this database, return an absolute URI.
MakeRelative (BSTR i_AbsoluteURI, BSTR * o_pRelativeURI);	Given an absolute URI within this database, return a relative URI within the scope of this database.
OpenDatabasePath (BSTR I_URI, BSTR * o_pDatabasePath);	Returns the directory path portion of a URI.
OpenDocumentName (BSTR i_URI, BSTR * o_pDocumentName);	Returns the document name portion of a URI.

OpenPersistRootPath (BSTR * o_pPath);	Returns the directory path to the root of the Groove persistent data directories.
OpenServiceName (BSTR i_URI, BSTR * o_pServiceName);	Returns the storage service portion of a URI.
Parse (BSTR i_URI, BSTR * o_pServiceName, BSTR * o_pDatabasePath, BSTR * o_pDocumentName);	Parses the pieces of the given URI.

Figure 10 illustrates the notification system interfaces. Interface 1000 (IGrooveLinkCallback) is an interface for use by a client of a storage manager that needs to be notified during the input processing of XML document or element when a definition for a link is found. The interface includes the methods defined in Table 5.

TABLE 5

Interface IGrooveLinkCallback : IDispatch

HandleLink (IGrooveElement * i_pLinkElement, IGrooveByteInputStream * i_pLinkData);	Called when the specified element contains a link attribute definition.
---	--

Another interface 1002 (IGrooveRPCServerCallback) is used by a client of a storage manager that needs to handle remote procedure calls (RPCs) on elements within XML documents. RPC server callbacks are a sub-class of the “util” base class (described below), that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveRPCServerCallback. Table 6 defines the methods used in the storage manager RPC server callback interface.

TABLE 6

Interface IGrooveElementRPCServerCallback : IDispatch

HandleCall (IGrooveElement * i_pInput,	Handle a RPC, receiving input parameters in
IgrooveElement ** o_ppOutput);	the Input element and returning output
	parameters in the Output element.

Figures 11,12 and 13 illustrate the document manipulation interfaces and the queue and RPC interfaces. In particular, Figure 11 shows the interfaces used to
 5 manipulate databases. An interface 1100 (IGrooveDatabase) is used by a client of a storage manager that needs to manage the databases in which documents are stored. It includes the methods in Table 7.

TABLE 7

Interface IGrooveDatabase : IDispatch

Checkpoint ();	Creates a durable point of state for the database.
ClearDataLost ();	Clears the database flag that indicates data may have been lost since the database was opened or the last transaction was committed.
CreateBinaryDocumentFromStream (IgrooveByteInputStream *i_pStream, BSTR I_DocumentName, IgrooveBinaryDocument ** o_ppDocument);	Creates a binary document with the specified name in the database.

CreateOrOpenXMLDocument (BSTR
i_DocumentName, BSTR
i_RootElementName, BSTR
i_SchemaURI, IUnknown *
i_pAdditionalSchemaURIs,
VARIANT_BOOL * o_pCreated,
IGrooveXMLDocument **
o_ppDocument);

Opens the specified XML document; creates an empty document with the specified name and schema if it doesn't already exist.

CreateXMLDocument (BSTR
i_DocumentName, BSTR
i_RootElementName, BSTR
i_SchemaURI, IUnknown *
i_pAdditionalSchemaURIs,
IGrooveXMLDocument **
o_ppDocument);

Creates an empty XML document with the specified name and schema in the database.

CreateXMLDocumentFromStream
(IGrooveByteInputStream * i_pStream,
GrooveParseOptions i_ParseOptions,
BSTR i_DocumentName, BSTR
i_SchemaURI, IUnknown *
i_pAdditionalSchemaURIs, IUnknown *
i_pLinkCallback,
IGrooveXMLDocument **
o_ppDocument);

Given a stream of bytes, representing one of the supported character set encodings of a XML document, creates an XML document in the database.

DeleteDocument (BSTR
i_DocumentName);

Deletes the named document.

DocumentExists (BSTR
i_DocumentName, VARIANT_BOOL *
o_pDocumentExists);

Given the specified document name, checks for the existence of the document in the database.

IsTransactionInProgress
(VARIANT_BOOL *
o_pTransactionInProgress);

Returns TRUE if a transaction is in progress.

OpenBinaryDocument (BSTR
i_DocumentName,
IGrooveBinaryDocument **
o_ppDocument);

Opens the specified binary document.

OpenCrossProcessSemaphore (BSTR
i_Name, VARIANT_BOOL
i_Reentrant,
IGrooveCrossProcessSemaphore **
o_ppSemaphore);

Creates a new cross process synchronization object. If Name is not specified, the default name for the database is used. If the semaphore is not Reentrant, repeated attempts to lock the semaphore within the same thread and process will block.

OpenDocumentNameEnum
(VARIANT_BOOL i_OpenOnly,
IGrooveBSTREnum **
o_ppDocumentNames);

Returns an enumeration of the documents currently in a database.

OpenTransaction (VARIANT_BOOL i_BeginLock, VARIANT_BOOL i_ReadOnly, VARIANT_BOOL i_BeginTransaction, VARIANT_BOOL i_Reentrant, BSTR i_LockName, IGrooveTransaction ** o_ppTransaction);	Creates a new transaction on the database. BeginLock specifies whether the database cross process semaphore should be locked. BeginTransaction specifies whether the transaction should start now. If LockName is not specified, the default name for the database is used. If the semaphore is not Reentrant, repeated attempts to lock the semaphore within the same thread and process will block.
OpenURI (BSTR * o_pDatabaseURI);	Returns the URI for this database.
OpenXMLDocument (BSTR i_DocumentName, IGrooveXMLDocument ** o_ppDocument);	Opens the specified XML document.
WasDataLost (VARIANT_BOOL * o_pDataLost);	Returns the value of a flag indicating whether data may have been lost since the database was opened or the last transaction was committed.

Table 8 illustrates the methods for an interface 1102 (IGrooveCrossProcessSemaphore) for a client of a storage manager that needs to synchronize access among processes.

TABLE 8

Interface IGrooveCrossProcessSemaphore : IDispatch

DoLock (VARIANT_BOOL i_ReadOnly);	Locks the semaphore. If ReadOnly is TRUE, only retrieval operations may be performed on the database, otherwise, any operation may be performed.
DoUnlock ();	Unlocks the semaphore.

Table 9 illustrates an interface 1104 (IGrooveTransaction) for a client of a storage manager that needs to group operations within a database. Transactions are a sub-class of cross-process semaphores, that is, all of the methods for IGrooveCrossProcessSemaphore also apply to IGrooveTransaction. The storage manager transaction interface includes the following methods:

TABLE 9

Interface IGrooveTransaction : IGrooveCrossProcessSemaphore

Abort ();	Ends the transaction. All work done to the database since the start of the transaction is discarded.
Begin (VARIANT_BOOL i_ReadOnly);	Starts a transaction. If ReadOnly is false, the database may be updated.
BeginIndependent (VARIANT_BOOL i_ReadOnly);	Starts another transaction for this thread. Only one independent transaction is allowed per thread.
Commit ();	Ends the transaction. All work done to the database since the start of the transaction is reliably stored in the database.

Figure 12 shows interfaces which allows clients of the storage manager to manipulate documents and elements within those documents. Table 10 illustrates an interface 1200 (IGrooveDocument) for a client of a storage manager that needs to manage documents within a database. The storage manager document interface includes the following methods:

TABLE 10

Interface IGrooveDocument : IDispatch

OpenCrossProcessSemaphore (BSTR i_Name, VARIANT_BOOL i_Reentrant, IgrooveCrossProcessSemaphore ** o_ppSemaphore);	Creates a new cross process synchronization object. If Name is not specified, the URI for the document is used. If the semaphore is not Reentrant, repeated attempts to lock the semaphore within the same thread and process will block.
OpenDatabase (IGrooveDatabase ** o_ppDatabase);	Returns an interface to the database object that contains this document.
OpenName (BSTR * o_pDocumentName);	Returns the document name.
OpenURI (BSTR * o_pURI);	Returns the URI that identifies this document.

Table 11 illustrates an interface 1202 (IGrooveXMLDocument) for a client of a storage manager that needs to manage XML documents within a database. XML documents are a sub-class of documents, that is, all of the methods for IGrooveDocument also apply to IGrooveXMLDocument. The storage manager XML document interface includes the following methods:

TABLE 11

interface IGrooveXMLDocument : IGrooveDocument

GenerateGrooveID (BSTR i_GrooveIDBase, double * o_pGrooveID);	Generates an 8 byte identifier from the string identifier i_GrooveIDBase.
ConvertGrooveIDToSerializedGrooveID (double i_GrooveID, BSTR * o_pGrooveIDString);	Converts an 8 byte identifier to the string i_GrooveID.
ConvertSerializedGrooveIDToGrooveID (BSTR i_GrooveIDString, double * o_pGrooveID);	Converts a string version of a Groove identifier to an 8 byte version.
CreateElement (BSTR i_Name, IUnknown * i_pParent, IGrooveElement ** o_ppElement);	Creates a new element with the supplied Tag; the tag cannot be altered once created. If a Parent reference is supplied, the new element is created as a child of that parent.
CreateElementCopy (IGrooveElement * i_pSource, IGrooveElement * i_pParent, VARIANT_BOOL i_ShallowCopy, IGrooveElement ** o_ppElement);	Does a deep/shallow copy of the specified element and all of its children (recursively for deep; just the one level for shallow), putting the new element(s) in under the Parent element.
CreateElementFromSchema (BSTR i_Name, IGrooveElement * i_pParent, IGrooveElement ** o_ppElement);	Creates an element that conforms to the element's definition in the schema. Creates the element, its attributes, and any child elements.

CreateElementFromStream

(IGrooveByteInputStream * i_pStream,
GrooveParseOptions i_ParseOptions,
IUnknown * i_pParent, IUnknown *
i_pLinkCallback, IGrooveElement **
o_ppElement);

Using a parser, creates an element, reads from a byte input stream and creates elements and attributes from the text stream as necessary, inserting them into the element, which is then returned to the caller. If a Parent reference is supplied, the new element is created as a child of that parent.

CreateLocator (IGrooveLocator **
o_ppLocator);

Returns the interface to a new locator object.

FindElementByID (BSTR i_ID,
IGrooveElement ** o_ppElement,
VARIANT_BOOL * o_pFound);

Looks for an element of the specified ID and returns a boolean value if found.

OpenElementByID (BSTR i_ID,
IGrooveElement **o_ppElement);

Looks for an element of the specified ID.

OpenElementEnumByAttributeValue
(BSTR i_ElementName, BSTR
i_AttributeName, BSTR
i_AttributeValue, IGrooveElementEnum
**o_ppElementEnum);

Returns an enumeration of all of the elements within the document that have the named attribute with the specified value.

OpenElementEnumByAttributeValueAs
Bool (BSTR i_ElementName, BSTR
i_AttributeName, VARIANT_BOOL
i_AttributeValue, IGrooveElementEnum
**o_ppElementEnum);

Returns an enumeration of all of the elements within the document that have the named attribute with the specified boolean type value.

OpenElementEnumByAttributeValueAs Double (BSTR i_ElementName, BSTR i_AttributeName, double i_AttributeValue, IGrooveElementEnum **o_ppElementEnum);	Returns an enumeration of all of the elements within the document that have the named attribute with the specified double floating type value.
OpenElementEnumByAttributeValueAs Long (BSTR i_AttributeName, long i_AttributeValue, IGrooveElementEnum **o_ppElementEnum);	Returns an enumeration of all of the elements within the document that have the named attribute with the specified long integer type value.
OpenElementEnumByLocator (BSTR i_LocatorText, IGrooveElementEnum ** o_ppElementEnum);	Returns an element enumerator with references to all elements satisfying the specified element locator expression. If there are no matching elements, the element enumerator will be created with no contents.
OpenElementEnumByName (BSTR i_Name, IGrooveElementEnum ** o_ppElementEnum);	Returns an enumeration of all of the elements within the document that have the specified tag name.
OpenMetaElement (IGrooveElement ** o_ppElement);	Returns the interface to the meta element that defines this XML document.
OpenRootElement (IGrooveElement ** o_ppRootElement);	Opens the root element for the XML document.

Table 12 illustrates the methods for an interface 1204 (IGrooveBinaryDocument) for a client of a storage manager that needs to manage binary documents within a database. Binary documents are a sub-class of documents, that is, all of the methods for IGrooveDocument also apply to IGrooveBinaryDocument.

TABLE 12

interface IGrooveBinaryDocument : IGrooveDocument

OpenByteInputStream (IGrooveByteInputStream ** o_ppByteInputStream);	Returns the interface to a byte stream object that can be used to read bytes within the binary document.
--	--

Table 13 illustrates an interface 1206 (IGrooveLocator) for a client of a storage manager that needs to search for elements using locator queries as defined in a specification called XSLT. Details of the XSLT specification can be found at <http://www.w3.org/TR/xslt>. The storage manager locator interface includes the following methods:

TABLE 13

interface IGrooveLocator : IDispatch

FindElement (BSTR i_LocatorStr, IGrooveElement * i_pContextElement, IGrooveElement ** o_ppElement, VARIANT_BOOL * o_pFound);	Returns an interface to the element object that satisfies the search specified by the Locator string within the scope of the context element.
Invalidate (VARIANT_BOOL i_AssignNewIDs);	Clears the state information in the interface instance.
OpenElementEnum (BSTR i_LocatorStr, IGrooveElement * i_pContextElement, VARIANT_BOOL i_Sort, BSTR i_SortConstraint, BSTR i_SortKey, GrooveSortOrder i_SortOrder, IGrooveElementEnum ** o_ppElements);	Returns an enumerator of all elements that match the Locator string, collated according to the specified sorting criteria.

OpenElementEnumWithTumblers (BSTR i_LocatorStr , IGrooveElement *i_pContextElement, VARIANT_BOOL i_RelativeTumblers, IGrooveBSTREnum ** o_ppTumblers, VARIANT_BOOL i_Sort, BSTR i_SortConstraint, BSTR i_SortKey, GrooveSortOrder i_SortOrder, IGrooveElementEnum ** o_ppElements);	Perform the search specified by the Locator string on the elements pointed to by the context element, returning the tumbler values for each match as well as the matching elements, collated according to the specified sorting criteria.
OpenText (BSTR i_LocatorStr, IGrooveElement * i_pContextElement, BSTR * o_pValue);	Returns the text from element or attribute that satisfies the search specified by the Locator string within the scope of the context element.

Table 14 illustrates an interface 1208 (IGrooveTransform) for a client of a storage manager that needs to perform XML document transformations as defined in XSLT. The storage manager transform interface includes the following methods:

TABLE 14

Interface IGrooveTransform : IDispatch

TransformXMLDocument (IGrooveXMLDocument * i_pXMLDocument, IGrooveElement * i_pStartElement, BSTR i_SortRule, long i_StartElementNum, long i_NumElements, IGrooveXMLDocument * io_pResultDocument, VARIANT_BOOL i_AlwaysOutputHeader, long * o_pElementsProcessed);	Transforms the input XML document, returning the result of the transformation in ResultDocument.
TransformElement (IGrooveElement * i_pContextElement, BSTR i_TransformationTemplate, IGrooveXMLDocument ** o_ppResultDocument);	Transforms the input ContextElement, returning the result of the transformation in ResultDocument.

Table 15 illustrates an interface 1210 (IGrooveElement) which allows a client of a storage manager to manipulate elements within XML documents. The storage manager element interface includes the following methods:

TABLE 15

Interface IGrooveElement : IDispatch

AppendContent (BSTR i_Text, GrooveContentType i_Type);	Inserts the kind of content as the last of its type within this element.
---	---

AppendContentElement
(IGrooveElement * i_pElement);

Inserts the element as the last content element.

AppendContentProcessingInstruction
(BSTR i_Target, BSTR i_Text);

Inserts a processing instruction, with target Target, as the last processing instruction.

CreateElement (BSTR i_Name,
IGrooveElement * i_pParent,
IGrooveElement ** o_ppElement);

Create a new element in the same document.

CreateElementCopy (IGrooveElement *
i_pSource, IGrooveElement * i_pParent,
VARIANT_BOOL i_ShallowCopy,
IGrooveElement ** o_ppElement);

Does a deep/shallow copy of the specified element and all of its children (recursively for deep; just the one level for shallow), putting the new element(s) in the destination document. The returned element must be attached into the document's element tree.

CreateElementFromSchema (BSTR
i_Name, IGrooveElement * i_pParent,
IGrooveElement ** o_ppElement);

Creates an element that conforms to the element's definition in the schema. Creates the element, its attributes, and any child elements.

CreateElementRPCClient
(IGrooveElementRPCClient
**o_ppRPCClient);

Creates and returns the interface to the element RPC client.

CreateElementRPCServer
(IGrooveElementRPCServer **
o_ppRPCServer);

Creates and returns the interface to the element RPC server.

CreateElementRPCServerThread
 (IGrooveElementRPCServerCallback *
 i_pCallback,
 IGrooveElementRPCServerThread **
 o_ppRPCServerThread);

Creates and returns the interface to the
 element RPC server thread.

CreateLink (IGrooveDocument *
 i_pDocument, BSTR i_Title, BSTR
 i_Role, GrooveXLinkShow i_Show,
 GrooveXLinkActuate i_Actuate,
 GrooveXLinkSerialize i_Serialize);

Creates a link to another document, using
 the specified XLink parameters.

DecrementAttributeAsLong (BSTR
 i_Name, long * o_pOldValue);

Subtracts 1 from the value of a long integer
 type attribute.

Delete ();

Permanently removes the element from the
 document. No further operations may be
 performed on a deleted element

DeleteAllAttributes ();

Removes all attributes from the element.

DeleteAllContent ();

Removes all child content elements and text
 from the element and deletes them from the
 document.

DeleteAttribute (BSTR i_Name);

Removes the named attribute from the
 element.

DeleteContent (long i_Ordinal);

Removes the content at the specified
 position from the element.

DeleteLinkAttributes ();

Removes all attributes that are links from the element.

DetachFromParent ();

Removes this element from the content of its parent. The element is still part of the document and must be reattached or destroyed before it is released.

DoesAttributeExist (BSTR i_Name,
VARIANT_BOOL * o_pFound);

Returns whether the attribute is set on the element.

Duplicate (IGrooveElement *
i_pTargetElement, VARIANT_BOOL
i_ShallowDuplicate);

Make the specified target element a duplicate of this element, overriding attributes and, if ShallowDuplicate is FALSE, all descendent elements.

FindAttribute (BSTR i_Name, BSTR *
o_pValue, VARIANT_BOOL *
o_pFound);

Gets any arbitrary attribute as text. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsBinary (BSTR i_Name,
IGrooveByteInputStream ** o_ppValue,
VARIANT_BOOL *o_pFound);

Gets any arbitrary attribute as Binary. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsBinaryArray (BSTR
i_Name, SAFEARRAY(BYTE) *
o_ppValue, VARIANT_BOOL *
o_pFound);

Gets any arbitrary attribute as Binary and return the value in an array. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsBinaryToStream (BSTR
i_Name, IGrooveByteOutputStream *
i_pStream, VARIANT_BOOL
*o_pFound);

Gets any arbitrary attribute as Binary and returns the value in a stream. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsBool (BSTR i_Name,
VARIANT_BOOL * o_pValue,
VARIANT_BOOL * o_pFound);

Gets any arbitrary attribute as Boolean. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsDouble (BSTR i_Name,
double * o_pValue, VARIANT_BOOL *
o_pFound);

Gets any arbitrary attribute as Double. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsGrooveID (BSTR
i_Name, double * o_pValue,
VARIANT_BOOL * o_pFound);

Gets any arbitrary attribute as a Groove identifier. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsLong (BSTR i_Name,
long * o_pValue, VARIANT_BOOL *
o_pFound);

Gets any arbitrary attribute as Long. The attribute must have been set as the given type or be specified as that type in the document schema. If the attribute is not in the element, Found is FALSE and no value is returned.

FindAttributeAsVARIANT (BSTR
i_Name, VARIANT * o_pValue,
VARIANT_BOOL * o_pFound);

Gets any arbitrary attribute as a variant value. If the attribute is not in the element, Found is FALSE and no value is returned.

FindContentElementByName (BSTR
i_Name, IGrooveElement **
o_ppElement, VARIANT_BOOL *
o_pFound);

Within the context of this element, find an element with the specified tag name. If the element is not found, Found is FALSE and no element reference is returned.

FindContentElementByNameAndAttribute (BSTR i_Name, BSTR
i_AttributeName, BSTR i_AttributeValue,
IGrooveElement ** o_ppElement,
VARIANT_BOOL * o_pFound);

Within the context of this element, find an element with the specified tag name and attribute name with the specified attribute value. If the element is not found, Found is FALSE and no element reference is returned

FindParent (IGrooveElement **
o_ppParent, VARIANT_BOOL *
o_pFound);

Gets an object's parent element. An element can have only a single parent and may only be referenced from a single content entry of a single element. If the element does not have a parent, Found is FALSE and no value is returned.

GetActuate (GrooveXLinkActuate *
o_pActuate);

Returns the value of the Actuate parameter in this element's link attribute.

GetAttributeCount (long * o_pCount);

Returns the number of attributes an element has.

GetContentCount (long * o_pCount);

Returns the number of content and text entries in this element.

GetContentType (long i_Ordinal,
GrooveContentType * o_pType);
GetOrdinal (long * o_pOrdinal);

Returns the type of content at the specified ordinal position.

Gets the ordinal position within the parent's content of this element.

GetSerialize (GrooveXLinkSerialize *
o_pSerialize);

Returns the value of the Serialize parameter in this element's link attribute.

GetShow (GrooveXLinkShow *
o_pShow);

Returns the value of the Show parameter in this element's link attribute.

IncrementAttributeAsLong (BSTR
i_Name, long * o_pOldValue);

Adds 1 to the value of a long integer type attribute.

InsertContent (long i_Ordinal, BSTR i_Text, GrooveContentType i_Type);	Inserts the text entry at the specified ordinal location
InsertContentElement (long i_Ordinal, IGrooveElement * i_pElement);	Inserts the element at the specified ordinal location
InsertContentProcessingInstruction (long i_Ordinal, BSTR i_Target, BSTR i_Text);	Inserts a Text processing instruction, with target Target, at the specified ordinal position.
IsLinkElement (VARIANT_BOOL * o_pIsLink);	Determines whether or not the element contains XLink markup.
IsReferenced (VARIANT_BOOL * o_pIsReferenced);	Returns TRUE if this element is referenced.
IsSame (IGrooveElement * i_pElement, VARIANT_BOOL * o_pIsSame);	Returns TRUE if the specified element object is this element or equal to this element.
OpenAttribute (BSTR i_Name, BSTR *o_pValue);	Gets any arbitrary attribute as text.
OpenAttributeAsBinary (BSTR i_Name, IGrooveByteInputStream ** o_ppValue);	Gets any arbitrary attribute as Binary. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsBinaryArray (BSTR
i_Name, SAFEARRAY(BYTE) *
o_ppValue);

Gets any arbitrary attribute as Binary and return the value in an array. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsBinaryToStream (BSTR
i_Name, IGrooveByteOutputStream *
i_pStream);

Gets any arbitrary attribute as Binary and returns the value in a stream. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsBool (BSTR i_Name,
VARIANT_BOOL * o_pValue);

Gets any arbitrary attribute as Boolean. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsDouble (BSTR i_Name,
double * o_pValue);

Gets any arbitrary attribute as Double. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsGrooveID (BSTR
i_Name, double * o_pValue);

Gets any arbitrary attribute as a Groove identifier. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsLong (BSTR i_Name,
long * o_pValue);

Gets any arbitrary attribute as Long. The attribute must have been set as the given type or be specified as that type in the document schema.

OpenAttributeAsVARIANT (BSTR
i_Name, VARIANT * o_pValue);

Gets any arbitrary attribute as a variant value.

OpenAttributeEnum
(IGrooveStringStringEnum **
o_ppAttributes);

Enumerates all of the element's attributes as text.

OpenAttributeVariantEnum
(IGrooveNameValueEnum **
o_ppEnum);

Enumerates all of the element's attributes as variant data types.

OpenBoundCode (IGrooveBoundCode
** o_ppBoundCode);

Returns an instance of the object bound to the element.

OpenContentComment (long i_Ordinal,
BSTR * o_pComment);

Returns the text of the comment that is a contained in this element at the specified Ordinal position.

OpenContentElement (long i_Ordinal,
IGrooveElement ** o_ppElement);

Returns the child element interface that is a contained in this element at the specified Ordinal position.

OpenContentElementByName (BSTR
i_Name, IGrooveElement **
o_ppElement);

Within the context of this element, find an element with the specified tag name and return its interface.

OpenContentElementByNameAndAttribute (BSTR i_Name, BSTR i_AttributeName, BSTR i_AttributeValue, IGrooveElement ** o_ppElement);	Within the context of this element, find an element with the specified tag name and attribute name with the specified attribute value.
OpenContentElementEnum (IGrooveElementEnum ** o_ppElements);	Returns an enumeration of all child content elements (non-recursively).
OpenContentElementEnumByName (BSTR i_Name, IGrooveElementEnum ** o_ppElements);	Returns an enumeration of all child content elements (non-recursively). Only elements with the given name will be returned.
OpenContentElementEnumByNameAndAttribute (BSTR i_Name, BSTR i_AttributeName, BSTR i_AttributeValue, IGrooveElementEnum ** o_ppElements);	Returns an enumeration of all content elements within the scope of this element that have the specified tag name and attribute name with the specified attribute value.
OpenContentProcessingInstruction (long i_Ordinal, BSTR * o_pTarget, BSTR * o_pText);	Returns the XML processing instruction at the specified ordinal position.
OpenContentProcessingInstructionTarget (long i_Ordinal, BSTR * o_pTarget);	Returns the target of the XML processing instruction at the specified ordinal position.
OpenContentProcessingInstructionText (long i_Ordinal, BSTR * o_pText);	Returns the PI text of the XML processing instruction at the specified ordinal position.
OpenContentText (long i_Ordinal, BSTR * o_pText);	Returns the context text at the specified ordinal position.

OpenContentTextEnum

(IGrooveBSTREnum ** o_ppText);

OpenElementQueue

(IGrooveElementQueue ** o_ppQueue);

Enumerates the text entries

(non-recursively).

Create an element queue on the element.

The element queue does not affect the element's structure.

OpenElementReferenceQueue

(IGrooveElementReferenceQueue **
o_ppQueue);

Returns the interface to reference queue object.

OpenHref (BSTR * o_pHref);

Returns the value of the HREF parameter in this element's link attribute.

OpenLinkAttributes (BSTR * o_pHref,
BSTR * o_pTitle, BSTR * o_pRole,
GrooveXLinkShow * o_pShow,
GrooveXLinkActuate * o_pActuate,
GrooveXLinkSerialize * o_pSerialize);

Retrieves all the standard link elements.

Note : not all the attributes are mandatory

OpenLinkedBinaryDocument
(VARIANT_BOOL i_SingleProcess,
IUnknown * i_pSecurityContext,
IGrooveBinaryDocument **
o_ppDocument);

Returns the interface to the binary document that is referenced in the HREF parameter in this element's link attribute.

OpenLinkedXMLDocument
(VARIANT_BOOL i_SingleProcess,
IUnknown * i_pSecurityContext,
IGrooveXMLDocument **
o_ppDocument);

Returns the interface to the XML document that is referenced in the HREF parameter in this element's link attribute.

OpenMultiReaderElementQueueReader (IGrooveMultiReaderElementQueueReader ** o_ppQueue); OpenMultiReaderElementQueueWriter (GrooveMultiReaderQueueOptions i_Options, IGrooveMultiReaderElementQueueWriter ** o_ppQueue);	Create an element multi-reader queue on the element and add a reader. This could change the structure of the element. Create an element multi-writer queue on the element and add a writer. This could change the structure of the element.
---	--

OpenMultiReaderElementReferenceQueueReader (IGrooveMultiReaderElementQueueReader ** o_ppQueue);	Returns the interface to the multi-reader element reference queue reader object.
--	--

OpenMultiReaderElementReferenceQueueWriter (GrooveMultiReaderQueueOptions i_Options, IGrooveMultiReaderElementQueueWriter ** o_ppQueue);	Returns the interface to the multi-reader element reference queue writer object.
---	--

OpenName (BSTR * o_pName);	Returns the element's tag name.
----------------------------	---------------------------------

OpenParent (IGrooveElement ** o_ppParent);	Gets an object's parent element. An element can have only a single parent and may only be referenced from a single content entry of a single element.
--	---

OpenReadOnlyElement (VARIANT_BOOL i_AllowOpenParent, IGrooveReadOnlyElement ** o_ppReadOnlyElement);	Return the read-only element interface to this element.
OpenReference (IGrooveElementReference ** o_ppElementReference);	Returns the element reference interface to this element.
OpenRole (BSTR * o_pRole);	Returns the value of the Role parameter in this element's link attribute.
OpenTitle (BSTR * o_pTitle);	Returns the value of the Title parameter in this element's link attribute.
OpenURI (BSTR * o_pName);	Returns the URI to this element.
OpenXMLDocument (IGrooveXMLDocument ** o_ppDocument);	Returns the interface pointer to the XML document containing this element.
Serialize (GrooveSerializeType i_Type, enum GrooveCharEncoding i_Encoding, GrooveSerializeOptions i_Options, IGrooveByteInputStream ** o_ppStream);	Serializes the element to a stream with the specified encoding and options.
SerializeReturnAdditionalLinkedDocume nts (GrooveSerializeType i_Type, enum GrooveCharEncoding i_Encoding, GrooveSerializeOptions i_Options, IGrooveDocumentEnum **	Serializes the element to a stream with the specified encoding and options. Returns an enumeration of interfaces to documents referenced by links in this element and all descendents.

```
o_ppAdditionalLinkedDocuments,
IGrooveByteInputStream **
o_ppStream);
SerializeToStream
(IGrooveByteOutputStream * i_pStream,
GrooveSerializeType i_Type, enum
GrooveCharEncoding i_Encoding,
GrooveSerializeOptions i_Options);
```

Serializes the element to a stream with the specified encoding and options.

```
SerializeToStreamReturnAdditionalLinke
dDocuments (IGrooveByteOutputStream
* i_pStream, GrooveSerializeType
i_Type, enum GrooveCharEncoding
i_Encoding, GrooveSerializeOptions
i_Options, IGrooveDocumentEnum **
o_ppAdditionalLinkedDocuments);
```

Serializes the element to a stream with the specified encoding and options. Returns an enumeration of interfaces to documents referenced by links in this element and all descendents.

```
SetAttribute (BSTR i_Name, BSTR
i_Value);
```

Sets any arbitrary attribute as text.

```
SetAttributeAsBinary (BSTR i_Name,
IGrooveByteInputStream * i_pValue);
```

Sets any arbitrary attribute as Binary. The attribute must have been set as the given type or be specified as that type in the document schema.

SetAttributeAsBinaryArray (BSTR
i_Name, SAFEARRAY(BYTE) *
i_pValue);

Sets any arbitrary attribute as Binary and returns the value in an array. The attribute must have been set as the given type or be specified as that type in the document schema.

SetAttributeAsBool (BSTR i_Name,
VARIANT_BOOL i_Value);

Sets any arbitrary attribute as Boolean. The attribute must have been set as the given type or be specified as that type in the document schema.

SetAttributeAsDouble (BSTR i_Name,
double i_Value);

Sets any arbitrary attribute as Double. The attribute must have been set as the given type or be specified as that type in the document schema.

SetAttributeAsGrooveID (BSTR i_Name,
double i_pValue);

Sets any arbitrary attribute as a Groove identifier. The attribute must have been set as the given type or be specified as that type in the document schema.

SetAttributeAsLong (BSTR i_Name, long
i_Value);

Sets any arbitrary attribute as Long. The attribute must have been set as the given type or be specified as that type in the document schema.

SetAttributeAsVARIANT (BSTR i_Name,
VARIANT * i_pValue);

Sets any arbitrary attribute using a Variant, which may be any variant type.

SetContent (long i_Ordinal, BSTR i_Text, GrooveContentType i_Type);	Sets the content as the type's ordinal position to the specified text. Note that content of different types have independent ordinal positions.
SetContentElement (long i_Ordinal, IGrooveElement * i_pElement);	Set the content element at the specified ordinal position.
SetContentProcessingInstruction (long i_Ordinal, BSTR i_Target, BSTR i_Text);	Set the content processing instruction at the specified ordinal position.
SetContentTextEnum (IGrooveBSTREnum * i_pEnum);	Creates text entries, separated by elements, for each text string in the enumerator.
SetLinkAttributes (BSTR i_Href, BSTR i_Title, BSTR i_Role, GrooveXLinkShow i_Show, GrooveXLinkActuate i_Actuate, GrooveXLinkSerialize i_Serialize);	Sets the link attributes needed to make the element a link element, including the 'xml:link' attribute, which is implicitly set to 'simple'.
SetName (BSTR i_Name);	Sets the name of the element.
SetTempAttribute (BSTR i_Name, BSTR i_Value);	Sets an attribute with a temporary value, which will not be committed in a transaction.

Table 16 illustrates the methods for an interface 1212 (IGrooveReadOnlyElement) for a client of a storage manager that needs to manipulate read-only elements within XML documents. Read-only elements are a sub-class of elements, that is, all of the methods for IGrooveElement also apply to IGrooveReadOnlyElement.

TABLE 16

interface IGrooveReadOnlyElement : IGrooveElement

OpenReadOnlyParent (IGrooveReadOnlyElement ** o_ppParent);	Returns a read-only element interface to the parent of this element.
OpenContentReadOnlyElement (long i_Ordinal, IGrooveReadOnlyElement ** o_ppElement);	Returns a read-only element interface to the content element at the specified Ordinal position.
OpenContentReadOnlyElementByName (BSTR i_Name, IGrooveReadOnlyElement ** o_ppElement);	Within the context of this element, find an element with the specified tag name and return its read-only interface.
FindContentReadOnlyElementByName (BSTR i_Name, IGrooveReadOnlyElement ** o_ppElement, VARIANT_BOOL * o_pFound);	Within the context of this element, find an element with the specified tag name and return its read-only interface. If the element is not found, Found is FALSE and no element reference is returned.
OpenContentReadOnlyElementEnum (IGrooveReadOnlyElementEnum ** o_ppElements);	Returns an enumeration of all child content elements read-only interfaces (non-recursively).
OpenContentReadOnlyElementEnumBy Name (BSTR i_Name, IGrooveReadOnlyElementEnum ** o_ppElements);	Returns an enumeration of all child content elements read-only interfaces (non-recursively). Only elements with the given name will be returned.

Table 17 illustrates an interface 1214 (IGrooveElementReference) for a client of a storage manager that needs to manipulate element references within XML documents. The storage manager element reference interface includes the following methods:

5

TABLE 17

Interface IGrooveElementReference : IDispatch

OpenElement	Returns a read-only element interface to the
(IgrooveReadOnlyElement **	referenced element.
o_ppElement);	

An interface 1216 (IGrooveElementUtilBase) for use within the storage manager's other interfaces is shown in Table 18. The IGrooveElementUtilBase is not an interface for commonly-used objects, but is intended to serve as the base class for other sub-classes (shown in Figure 13) that do have commonly-used objects. All of the "util" interfaces are associated with an element. The storage manager element util base interface includes the following methods:

TABLE 18

Interface IGrooveElementUtilBase : IDispatch

OpenDocument	Returns the interface of the containing XML
(IgrooveXMLDocument **	document.
o_ppDocument);	
OpenElement (IGrooveElement **	Returns the element's interface.
o_ppElement);	

Table 19 illustrates an interface 1218 (IGrooveBoundCode) for a client of a storage manager that needs to handle executable code associated with elements within

XML documents. The storage manager bound code interface includes the following methods:

TABLE 19

interface IGrooveBoundCode : IDispatch

SetElement (IGrooveElement * i_pElement);	Sets the element interface pointer associated with this element tag.
OpenElement (IGrooveElement ** o_ppElement);	Retrieves the element interface pointer associated with this element tag.

5

Figure 13 illustrates interfaces which are sub-classes of the
IGrooveElementUtilBase base class 1300, discussed above. Table 20 illustrates an
interface 1302 (IGrooveElementQueue) for a client of a storage manager that needs to
manipulate queues on elements within XML documents. Element queues are a sub-
class of the “util” base class, that is, all of the methods for IGrooveElementUtilBase also
apply to IGrooveElementQueue. The storage manager element queue interface
includes the following methods:

TABLE 20

interface IGrooveElementQueue : IGrooveElementUtilBase

Enqueue (IGrooveElement * i_pElement);	Enqueues the element. Note that the element must already be contained in the queue’s document.
---	--

Dequeue (long i_TimeoutMilliseconds, IGrooveElement ** o_ppElement);	Dequeues the next available element in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElement pointer will be NULL if the timeout period expires.
DequeueEnum (long i_TimeoutMilliseconds, IGrooveElementEnum ** o_ppElements);	Dequeues all available elements in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElement pointer will be NULL if the timeout period expires.
OpenEvent (IGrooveEvent ** o_ppEvent);	Returns an event that can be used to 'Wait' for an element to be enqueued

Table 21 illustrates an interface 1306 (IGrooveElementReferenceQueue) for a client of a storage manager that needs to manipulate queues on element references within XML documents. Element reference queues are a sub-class of the “util” base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveElementReferenceQueue. The storage manager element reference queue interface includes the following methods:

TABLE 21

interface IGrooveElementReferenceQueue : IGrooveElementUtilBase	
Enqueue (IGrooveElement * i_pElement);	Enqueues the element. Note that the element must already be contained in the queue's document.

EnqueueReference (IGrooveElement * i_pElement);	Enqueues a reference to the element. Note that the element must already be contained in the queue's document.
Dequeue (long i_TimeoutMilliseconds, IGrooveElementReference ** o_ppElementReference);	Dequeues the next available element in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElementReference pointer will be NULL if the timeout period expires.
DequeueEnum (long i_TimeoutMilliseconds, IGrooveElementReferenceEnum ** o_ppElementReferences);	Dequeues all available elements in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElementReferenceEnum pointer will be NULL if the timeout period expires.
OpenEvent (IGrooveEvent ** o_ppEvent);	Returns an event that can be used to 'Wait' for an element to be enqueued

Table 22 illustrates an interface 1310 (IGrooveMultiReaderElementQueueReader) for a client of a storage manager that needs to remove elements from multi-reader queues on elements within XML documents. Multi-reader element queues are a sub-class of the “util” base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveMultiReaderElementQueueReader. The storage manager multi-reader element queue reader interface includes the following methods:

TABLE 22

interface IGrooveMultiReaderElementQueueReader : IGrooveElementUtilBase

Dequeue (long i_TimeoutMilliseconds, IGrooveElement ** o_ppElement);	Dequeues the next available element in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElement pointer will be NULL if the timeout period expires.
DequeueEnum (long i_TimeoutMilliseconds, IGrooveElementEnum ** o_ppElements);	Dequeues all available elements in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElement pointer will be NULL if the timeout period expires.
OpenEvent (IGrooveEvent ** o_ppEvent);	Returns an event that can be used to 'Wait' for an element to be enqueued

Table 23 illustrates an interface 1314 (IGrooveMultiReaderElementQueueWriter) for a client of a storage manager that needs to add elements to multi-reader queues on elements within XML documents. Multi-reader element queues are a sub-class of the “util” base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveMultiReaderElementQueueWriter. The storage manager multi-reader element queue writer interface includes the following methods:

TABLE 23

interface IGrooveMultiReaderElementQueueWriter : IGrooveElementUtilBase

Enqueue (IGrooveElement *i_pElement, long * o_pNumEnqueued);	Enqueues the element and returns the number already enqueued. Note that the element must already be contained in the queue's document.
GetNumReaders (long * o_pNumReaders);	Get the number of readers on the queue.

Table 24 illustrates an interface 1318 (IGrooveMultiReaderElementReferenceQueueWriter) for a client of a storage manager that needs to add element references to multi-reader queues on elements within XML documents. Multi-reader element reference queues are a sub-class of the "util" base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveMultiReaderElementReferenceQueueWriter. The storage manager multi-reader element reference queue writer interface includes the following methods:

TABLE 24

**interface IGrooveMultiReaderElementReferenceQueueWriter :
IGrooveElementUtilBase**

Enqueue (IGrooveElement * i_pElement, long * o_pNumEnqueued);	Enqueues the element and returns the number already enqueued. Note that the element must already be contained in the queue's document.
--	---

EnqueueReference (IGrooveElement * i_pElement, long * o_pNumEnqueued);	Enqueues the element reference and returns the number already enqueued. Note that the element must already be contained in the queue's document.
GetNumReaders (long * o_pNumReaders);	Get the number of readers on the queue.

Table 25 illustrates an interface 1316 (IGrooveMultiReaderElementReferenceQueueReader) for a client of a storage manager that needs to remove element references from multi-reader queues on elements within XML documents. Multi-reader element reference queues are a sub-class of the “util” base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveMultiReaderElementReferenceQueueReader. The storage manager multi-reader element reference queue reader interface includes the following methods:

TABLE 25

interface IGrooveMultiReaderElementReferenceQueueReader :
IGrooveElementUtilBase

Dequeue (long i_TimeoutMilliseconds, IGrooveElementReference ** o_ppElementReference);	Dequeues the next available element reference in the queue. Returns only when an element is available or after the timeout period. The returned IGrooveElementReference pointer will be NULL if the timeout period expires.
DequeueEnum (long i_TimeoutMilliseconds, IGrooveElementReferenceEnum ** o_ppElementReferences);	Dequeues all available element references in the queue. Returns only when an element is available or after the timeout period. The returned

IGrooveElementReference pointer will be NULL if the timeout period expires.

OpenEvent (IGrooveEvent ** o_ppEvent); Returns an event that can be used to 'Wait' for an element to be enqueued

Table 26 illustrates an interface 1304 (IGrooveRPCClient) for a client of a storage manager that needs to perform remote procedure calls (RPCs) on elements within XML documents. RPC clients are a sub-class of the "util" base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveRPCClient. The storage manager RPC client interface includes the following methods:

TABLE 26

interface IGrooveElementRPCClient : IGrooveElementUtilBase

DoCall (IGrooveElement * i_pInput, IGrooveElement ** o_ppOutput); Make a RPC, using the Input element as the input parameters and receiving output parameters in the Output element.

SendCall (IGrooveElement * i_pInput); Make an asynchronous RPC, using the Input element as the input parameters.

OpenResponseQueue (IGrooveElementQueue ** o_ppQueue); Returns the queue where responses are received.

An interface 1308 (IGrooveRPCServerThread) for a client of a storage manager that needs to handle remote procedure calls (RPCs) on elements within XML documents is shown in Table 27. RPC server threads are a sub-class of the "util" base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveRPCServerThread. The storage manager RPC server callback interface has no

methods of its own, only those inherited from IGrooveElementUtilBase. It is provided as a distinct interface for type checking.

TABLE 27

interface IGrooveElementRPCServerThread : IGrooveElementUtilBase
(none)

Table 28 illustrates an interface 1312 (IGrooveRPCServer) for a client of a storage manager that needs to handle remote procedure calls (RPCs) on elements within XML documents. RPC servers are a sub-class of the “util” base class, that is, all of the methods for IGrooveElementUtilBase also apply to IGrooveRPCServer. The storage manager RPC server interface includes the following methods:

TABLE 28

interface IGrooveElementRPCServer : IGrooveElementUtilBase	
OpenCallQueue (IGrooveElementQueue ** o_ppQueue);	Returns the queue where calls are received.
SendResponse (IGrooveElement * i_pInput, IGrooveElement * i_pOutput, VARIANT_BOOL * o_bResult);	Sends a response to the caller, returning output parameters in the Output element.

The following tables illustrate allowed values for the enumerated data types listed in the above interfaces. In particular, Table 29, illustrates allowed values for the GrooveSerializeType enumerated data type.

TABLE 29

GrooveSerializeType

GrooveSerializeAuto	On input, Groove will determine the correct format by examining the first few bytes of the input stream. On output, Groove will select a format based on the kind of document or element data.
GrooveSerializeMIME	Format is MHTML, as defined in RFC 2557.
GrooveSerializeXML	Format is XML. Note that binary documents are not supported with this format, but it may be a body type in MHTML.
GrooveSerializeWBXML	Format is WBXML. Note that binary documents are not supported with this format, but it may be a body type in MHTML.

Table 30 illustrates the allowed values for the GrooveSerializeOptions enumerated data type.

TABLE 30

GrooveSerializeOptions

GrooveSerializeDefault	Use default serialization options.
GrooveSerializeWithFormatting	Indent, with blanks, each level of child content elements beneath the parent element.
GrooveSerializeSortedAttrs	Output the attributes for each element in order of ascending attribute name.

GrooveSerializeNoFragmentWrapper	Output without the fragment wrapper for document fragments (elements).
GrooveSerializeNoNamespaceContraction	Output with fully expanded element and attribute names.
GrooveSerializeNoProlog	Output without the XML document prolog.
GrooveSerializeNoLinks	Output without linked documents.
GrooveSerializeNotMinimum	Don't spend as much local processor time as needed to ensure the resulting output is the minimum size.

Table 31 illustrates the allowed values for the GrooveParseOptions enumerated data type.

5

TABLE 31

GrooveParseOptions	
GrooveParseDefault	Use default parse options.
GrooveParseStripContentWhitespace	Remove all extraneous whitespace from element content.
GrooveParseNoFragment	Parse a fragment that doesn't have a fragment wrapper.
GrooveParseNoNamespaceExpansion	Parse the document, but don't expand namespaces to their fully qualified form.
GrooveParseNoLinks	Parse a document and skip the links.

Table 32 illustrates the allowed values for the GrooveContentType enumerated data type.

TABLE 32

GrooveContentType	
GrooveContentElement	Content is a child element.
GrooveContentText	Content is body text.
GrooveContentCDATASection	Content is a CDATA section.
GrooveContentProcessingInstruction	Content is a processing instruction.
GrooveContentComment	Content is a comment.

Table 33 illustrates the allowed values for the GrooveXLinkShow enumerated data type.

TABLE 33

GrooveXLinkShow	
GrooveXLinkShowNew	New.
GrooveXLinkShowParsed	Parsed.
GrooveXLinkShowReplace	Replace

Table 34 illustrates the allowed values for the GrooveXLinkActuate enumerated data type:

TABLE 34

GrooveXLinkActuate

GrooveXLinkActuateUser	User.
GrooveXLinkActuateAuto	Auto.

Table 35 illustrates the allowed values for the GrooveXLinkSerialize enumerated data type.

5

TABLE 35

GrooveXLinkSerialize

GrooveXLinkSerializeByValue	By value.
GrooveXLinkSerializeByReference	By reference.
GrooveXLinkSerializeIgnore	Ignore.

Table 36 illustrates the allowed values for the GrooveMultiReaderQueueOptions enumerated data type.

10

TABLE 36

GrooveMultiReaderQueueOptions

GrooveMRQDefault	Use default options.
GrooveMRQAllReceive	All readers receive each event notification.
GrooveMRQEnqueueIfNoReaders	Enqueue even if no reader is currently queued to receive the element.

The fundamental data model of the storage manager is XML. XML is a semi-structured, hierarchical, hyper-linked data model. Many real world problems are not well represented with such complex structures and are better represented in tabular form. For example, spreadsheets and relational databases provide simple, tabular interfaces. In accordance with one aspect of the invention, in order to simplify the representation, XML structures are mapped to a tabular display, generally called a “waffle”. The waffle represents a collection of data. This mapping is performed by the collection manager, a component of the storage manager.

Collections are defined by a collection descriptor, which is an XML document type description. Like a document schema, the collection descriptor is a special kind of document that is stored apart from the collection data itself. There are many sources of collection data, but the primary source of collection data is a software routine called a record set engine. Driven by user commands, the record set engine propagates a set of updates for a collection to the collection manager. Based on those updates, the collection manager updates index structures and may notify waffle users via the notification system. When a waffle user needs updated or new collection data, the waffle user will call the collection manager to return a new result array containing the updated data. The waffle user may also navigate within the collection using cursors.

The following list shows the XML DTD contents for a collection descriptor document:

```
<!ELEMENT Collection ANY>
<!ATTLIST Collection
    Name      CDATA          #REQUIRED
    Start     (record|index) "record" #REQUIRED
    Version   CDATA          #REQUIRED
    Location  CDATA          #IMPLIED
>

<!ELEMENT Level (Column|Sorting|Level)*>
<!ATTLIST Level
    Mapping   (Flatten|Direct)
    Links     (Embed|Traverse) "Traverse"
>
```

```

<!ELEMENT Column EMPTY>
<!--ATTLIST Column
    Source          CDATA          #REQUIRED
    Output          CDATA          #REQUIRED
    MultiValue      (OnlyFirst|MultiLine|Concatenate)"    OnlyFirst"
    MultiValueSeparator      CDATA          #IMPLIED ","
-->

<!ELEMENT Sorting SortDescription+>

<!--ELEMENT SortDescription Group?|SortColumn+|Interval?>
<!--ATTLIST SortDescription
    Name          CDATA          #REQUIRED
-->

<!--ELEMENT SortColumn EMPTY>
<!--ATTLIST SortColumn
    Source          CDATA          #REQUIRED
    Order          (Ascending|Descending) #REQUIRED
    DataType       CDATA          #REQUIRED
    Strength       (Primary|Secondary|Tertiary|Identical) "Identical"
    Decomposition  (None|Canonical|Full) "None"
-->

<!--ELEMENT Group Group?|GroupColumn+>
<!--ATTLIST Group
    Grouping (Unique|Units) #REQUIRED
    GroupUnits (Years|Months|Days|Hours)
    AtGroupBreak (None|Count|Total) "None"
    Order (Ascending|Descending) #REQUIRED
    Strength (Primary|Secondary|Tertiary|Identical) "Identical"
    Decomposition (None|Canonical|Full) "None"
-->

<!--ELEMENT GroupColumn EMPTY>
<!--ATTLIST GroupColumn
    Source          CDATA          #REQUIRED
-->

<!--ELEMENT Interval EMPTY>
<!--ATTLIST Interval
    Start          CDATA          #REQUIRED
    End            CDATA          #REQUIRED
-->

```

Every Collection has a name that is used to reference the collection. The Start attribute specifies how to find the “root” of the collection. A collection with a record root is just a set of records, whereas a collection that starts with an index is navigated through the index and then the set of records. An index may be a concordance or full-text. The optional Location attribute is a relative URL that identifies where in the root to actually begin.

A Level defines the contents of part of the output hierarchy. A level consists of the columns in the level, the ordering or grouping of records in the level, and definitions of sub-levels. A level is associated with records in the source record stream through the Mapping attribute. If the mapping is Direct, a level represents a single source record type. If the mapping is Flatten, the level contains a source record type and all descendants of that record. The Flatten mapping may only be specified on the only or lowest level in the collection. The Links attribute specifies how records with link attributes should be handled. If links are Traversed, the record will be output as a distinct level. If links are Embedded, the child record of the source record will appear as though it is part of the source record.

A Column defines the mapping between a source field and the output array column. The Source attribute is a XSLT path expression in the source records. The Result attribute is a name of the field in the result array. The MultiValue and MultiValueSeparator attributes define how multi-valued source values are returned in the result.

Every collection must have at least one defined order. The order can be sorted collation or multi-level grouping with aggregate functions.

The SortColumn element defines the collation characteristics within a SortDescription. The Source attribute defines the name of the output column to be sorted. The Order must be either Ascending or Descending. The Strength and Decomposition values are input parameters that have the same meaning as defined in Unicode.

The two kinds of grouping are by unique values and by units. When a collection is grouped by unique values, all records with the same GroupColumn values will be

together in the same group – breaks between groups will occur at the change of GroupColumn values. When a collection is grouped by units, all records with the same GroupColumn values, resolved to the value of GroupUnits, will be together in the same group. For example, if GroupUnits is “Days”, all records for a given day will be in the same group. If AtGroupBreak is specified, a synthetic row will be returned that contains the result of the aggregate function at each value or unit break value.

The GroupColumn identifies the result column to be grouped.

The Interval identifies the two fields in each record that define a range. The datatypes of the Start and End columns must be either numeric or datetime.

The following example shows a collection descriptor document for a simple document discussion record view with six collation orders:

```
<Collection Name="Main" Start="Record" Version="0,1,0,0">
  <Level Mapping="Flatten">
    <Column Source="Title" Output="Title"/>
    <Column Source="_Modified" Output="_Modified"/>
    <Column Source="_CreatedBy" Output="_CreatedBy"/>
    <Sorting>
      <SortDescription Name="ByAscModified">
        <SortColumn Source="_Modified" Order="Ascending"
          DataType="DateTime"/>
      </SortDescription>
      <SortDescription Name="ByDescModified">
        <SortColumn Source="_Modified"
          Order="Descending" DataType="DateTime"/>
      </SortDescription>
      <SortDescription Name="ByAscAuthor">
        <SortColumn Source="_CreatedBy"
          Order="Ascending" DataType="String"/>
      </SortDescription>
      <SortDescription Name="ByDescAuthor">
        <SortColumn Source="_CreatedBy"
          Order="Descending" DataType="String"/>
      </SortDescription>
      <SortDescription Name="ByAscTitle">
        <SortColumn Source="Title" Order="Ascending"
          DataType="String"/>
      </SortDescription>
      <SortDescription Name="ByOrdinal">
        <SortColumn Source="" Order="Ordinal">
```

```

        DataType="Long"/>
      </SortDescription>
    </Sorting>
  </Level>
5 </Collection>

```

The following example shows a collection descriptor for a calendar view. Note the similarity to the prior example, but with a small change to the sort description, the collection is ordered by ranges of date intervals.

```

10 <Collection Name="Main" Start="Record" Version="0,1,0,0">
    <Level Mapping="Flatten">
        <Column Source="from-attributes(Subject)"
            Output="Subject"/>
        <Column Source="from-attributes(Start)"
            Output="Start"/>
15 <Column Source="from-attributes(End)"
            Output="End"/>
        <Column Source="from-attributes(RecurrenceEnd)"
            Output="RecurrenceEnd"/>
20 <Column Source="from-attributes(IsAllDay)"
            Output="IsAllDay"/>
        <Column Source="from-attributes(IsRecurrent)"
            Output="IsRecurrent"/>
        <Sorting>
25 <SortDescription Name="DateRanges">
            <Interval Start="Start" End="End"/>
        </SortDescription>
        </Sorting>
    </Level>
30 </Collection>

```

As is the basic storage manager, the collection manager is implemented in an object-oriented environment. Accordingly, both the collection manager itself and all of the collection components including collections, waffles, cursors, result arrays and the record set engine are implemented as objects. These objects, their interface, the underlying structure and the API used to interface with the collection manager are illustrated in Figure 14. The API is described in more detail in connection with Figure 15. Referring to Figure 14, the collection manager provides shared access to collections, via the collection manipulation API 1402, but, in order to enable a full

programming model for client applications, additional communication and synchronization operations are provided, within the context of a collection. For example, a user can control a record set engine 1412 by means of the engine API 1404. Under control of commands in the engine API 1404, the record set engine 1412 propagates a set of updates for a collection to the distributed virtual object system 1410 that is discussed above. Based on those updates, the distributed virtual object system 1410 updates index and other structures.

Other client components may need to be aware of changes within components, such as waffles, managed by the collection manager. Accordingly, the collection manager provides an interface 1400 to an interest-based notification system 1406 for those client components. The notification system 1406 provides notifications to client component listeners who have registered an interest when values within objects 1408 that represent a collection change.

Collection data is represented by a set of objects including collection objects, record objects, waffle objects, cursor objects and result array objects 1408. The objects can be directly manipulated by means of the collection manipulation API 1402. The collection related objects 1408 are actually implemented by the distributed virtual object system 1410 that was discussed in detail above.

Figure 15 and the following tables comprise a description of the interfaces for each of the objects used to implement a preferred embodiment of the inventive collection manager. As with the storage manager implementation, these objects are designed in accordance with the Common Object Model (COM), but could also be implemented using other styles of interface and object model.

Table 37 illustrates an interface 1500 (IGrooveCollectionManager) for a collection manager that encapsulates the basic framework for the major operations performed on a collection. The collection manager interface includes the following methods:

TABLE 37

Interface IGrooveCollectionManager : IGrooveDispatch

CreateCollection(IGrooveElement *i_pCollectionDescriptor, BSTR i_CollectionURL, BSTR i_EngineID, IGrooveCollection **o_ppCollection);	Creates a new collection object. The CollectionDescriptor should contain a collection descriptor in XML according to the GrooveCollection XML DTD.
DeleteCollection(IGrooveXMLDocumen t *i_pSourceDocument, BSTR i_CollectionURL);	Deletes the specified collection from the SourceDocument.
OpenCollection(IGrooveElement *i_pCollectionDescriptor, BSTR i_CollectionURL, BSTR i_EngineID, IGrooveCollection **o_ppCollection);	Opens an existing collection object.
OpenCollectionEnum(IGrooveXMLDoc ument *i_pSourceDocument, IGrooveBSTREnum **o_ppCollectionNames);	Return an enumeration of all collections within a document.
ParseCollectionDescriptor(IGrooveEle ment *i_pCollectionElement, void * m_Levels);	Creates a collection document according to the specified collection descriptor.
UpdateCollection(void *i_Updates, BSTR i_EngineID, IGrooveElement ** o_ppUpdateContext);	Perform the requested sequence of operations (of kind GrooveCollectionUpdateOp) on the collection for EngineID.

Table 38 illustrates an interface 1502 (IGrooveCollection) for a collection that encapsulates the basic framework for the major operations performed on a collection. The collection interface includes the following methods:

5

TABLE 38

Interface IGrooveCollection : IGrooveDispatch

AdviseListeners(IGrooveElement *i_UpdateContext);	Notifies subscribing listeners of changes to this element.
CloseWaffle(IGrooveWaffle *i_pWaffle);	Removes an IGrooveWaffle instance from the list of the collection's listeners.
Delete(void);	Deletes the collection from the database.
DisableListeners (void);	Disables event notifications for all subscribing listeners.
EnableListeners (void);	Enables event notifications for all subscribing listeners. Event notifications are enabled by default, so this is only necessary if DisableListeners was previously called.
Find(BSTR i_pQuery, IGrooveCollection ** o_ppQueryResult);	Using the specified XSLT query expression, evaluate it on the collection and return a new collection as the result.

XSLT locators have the form:

AxisIdentifier(NodeTest Predicate)

where AxisIdentifier is one of:

from-ancestors

from-ancestors-or-self

from-attributes

from-children
from-descendants
from-descendants-or-self
from-following
from-following-siblings
from-parent
from-preceding
from-preceding-siblings
from-self
from-source-link

NodeTest is of the form QName and tests whether the

node is an element or attribute with the specified name.

A Predicate is of the form [PredicateExpr]

PredicateExpr is a Expr

Expr is one of:

VariableReference

(*Expr*)

Literal

Number

FunctionCall

Multiple predicates are separated by "/"

For example:

from-children(ElementName[from-attributes(AttributeName)])

GetCursor(IGrooveCollectionCurs
or **o_ppCursor); Returns a copy of the cursor currently used by the
collection.

GetCursorPosition(double *
o_pRelativePosition); Returns the relative position of the cursor as a
number between 0.0 (first row) and 100.0 (last
row).

GetEngineMappingTable(void
**o_ppEngineURLs); Returns the engine mapping table.

GetExpansionMask(long
*o_pMask); Gets the current value of the expansion mask.

GetRecordCount(long *
o_pRecordCount); Returns the number of records in the collection.

HasOrdinalSort(BSTR *
o_pSortName, VARIANT_BOOL
*o_pHaveSort); If the collection has an ordinal index, returns the
sort name and the value TRUE, otherwise it
returns FALSE.

HasSort(BSTR i_ColumnName,
GrooveCollationOrder
i_CollationOrder, long i_Level,
BSTR *o_pSortName,
VARIANT_BOOL *o_pHaveSort); Returns a bool indicating whether or not a sort
exists in the collection for the column specified by
i_ColumnName on level i_Level in collation order
i_AscendingSort. If a sort exists the sort name is
returned in o_pSortName.

IsEmpty(VARIANT_BOOL
*o_plsEmpty); Returns a bool indicating whether or not the
collection is empty.

MarkAll(VARIANT_BOOL i_Read); Sets the record read/unread indicator for all records in the collection to be the value of Read.

MarkRead(double i_RecordID); Sets a specific record to be marked as read.

MarkUnread(double i_RecordID); Sets a specific record to be marked as unread.

MoveCursor(GrooveCollectionCursorPosition i_AbsolutePosition, GrooveCollectionNavigationOp i_Navigator, long i_Distance, long *o_pDistanceMoved); Every collection has a cursor. The cursor establishes the starting position in the source document, which will then be used to build the result document. AbsolutePosition may have the values First, Last, or Current.

Navigator may have the following values:

Value

Description

NextAny, PriorAny

Move the cursor to the next/previous source row, traversing down through child rows and up through parent rows.

NextPeer, PriorPeer

Move the cursor to the next/previous source row at the same level, stopping if a row at a higher level is reached.

NextParent, PriorParent

Move the cursor to the next/previous parent source row, traversing until the root row is

reached.

NextData, PriorData

Move the cursor to the next/previous row that contains a data record.

NextUnread, PriorUnread

Move the cursor to the next/previous unread row.

Distance sets the numbers of iterations to move the cursor, starting at AbsolutePosition and moving through Distance iterations of Navigator movement.

MoveCursor returns the number of iterations the cursor was actually moved.

Sets the collection's cursor to point to the specified record.

MoveCursorToRecord(double
i_RecordID);

MoveCursorToValue(BSTR
i_pQuery, double * o_pRecordID);

Using the current sort order, positions the cursor to the row that meets the criteria of matching the relop to the input query values. The relop (relational operator) may be EQ, LT, LE, GT, or GE. The query values must match, in order, the datatypes of the columns of the current sort order or must be able to be converted in a loss-less manner to those datatypes. Fewer query values may be specified than are defined in the sort order, which will result in a partial match. For collections ordered on an interval, the first query value is the interval's starting value and the second is the ending value.

MoveToCursor(IGrooveCollection
Cursor *i_pCursor);

Moves the collection to the position specified by i_pCursor.

Open(BSTR i_CollectionURL,
IGrooveElement
*i_pCollectionDescriptorElement,
VARIANT_BOOL i_Temp,
VARIANT_BOOL i_Shared,
VARIANT_BOOL * o_pCreated);

Creates or opens the collection specified by I_CollectionURL within the Groove storage service i_ServiceType. Returns a bool indicating whether or not the collection was created for the first time.

OpenRecord(double i_RecordID,
IGrooveRecord ** o_ppRecord);

Returns an interface pointer to a specific record in the collection.

OpenRecordID(double
i_SourceRecordID, enum
GrooveCollectionNavigationOp
i_Relation, double *

Starting from the position of the SourceRecordID, perform the specified collection navigation operation and return the resulting record ID.

o_pTargetRecordID);

OpenResultArray(long
i_NumReturnRows, void
*io_pResultArray);

Given the collection's expansion mask, current cursor position and current sort order, return at most NumReturnRows into a result array conforming to the description below. Note that NumReturnRows is a quota only on the data rows – other synthesized header and footer rows may be returned as necessary.

Column Name

Data Type

Description

RowType

UINT1

==WAFFLE_ROW_DATA if the row is a data record returned from an engine,

==WAFFLE_ROW_HEADER false if the row is a synthesized header (e.g., category),

==WAFFLE_ROW_FOOTER if the row is a synthesized footer (e.g., aggregate result).

SynthKind

UINT1

If the row is a data row, this value is 0. If the row is a synthesized row, this value will be one of:

- BreakUnique: Indicates a change in value of categorized or sorted column. One of the ColumnName(i) columns will have the new

value.

- BreakUnitDay
- BreakUnitWeek
- BreakUnitMonth
- BreakUnitYear
- FuncTotal
- FuncCount

EngineID

UINT4

If the row is a data row: Index into the EngineID table, which is a vector of URLs stored as BSTRs.

If the row is a synthesized row, EngineID is 0.

RecordID

UINT4

If the row is a data row: RecordID returned from the engine identified by EngineID. RecordIDs are unique within EngineIDs.

If the row is a synthesized row: RecordID is a unique number within the collection.

Level

UINT1

Number of levels to indent this row. Level 0 is the top or outermost level.

RelativePosition

UINT2

A number between 0 and 10000 indicating the relative offset of this row from the beginning of the collection. [It may be an approximation.] For example, 6823 is the value for a row that is 68.23% of the way through the collection.

Read

BOOL

If the row is a data row: True if the [account??] has read the record. If the row is a synthesized row, Read is always true (even if it is collapsed).

ColumnName(i)

Defined by the collection descriptor.

Data value for this row/column. There will be as many columns in the array as there were defined columns at all levels.

OpenSchema(long i_Level,
VARIANT_BOOL
i_IncludeSystemColumns,
IGrooveRecordSchema
**o_ppCollectionSchema);

Return an interface pointer to the schema description for the records in the collection.

OpenTransaction(IGrooveTransact
ion **o_ppTransaction);

Creates a transaction on the collection document.

OpenWaffle(IGrooveWaffleListene
r *i_pListener, IGrooveWaffle

Creates an IGrooveWaffle instance and adds it to the collections list of event listeners .

****o_ppWaffle);**

**SetCursorPosition(double
i_RelativePosition);**

Sets the current position of the cursor to the row with the specified relative position. The position should be a number between 0.0 (first row) and 100.0 (last row).

SetExpansionMask(long i_Mask);

Sets the current value of the expansion mask. The mask is stored in a DWORD, but only the first 10 (or so) bits are used. If a bit is set, all data at the indicated level is expanded. The expansion mask is not persistent or shared – its effect is only on this collection object. The default value of the expansion mask is all 1s.

**SetRecordExpansion(double
i_RecordID, VARIANT_BOOL
i_Expand);**

Sets the expansion state for a single row for this scope. If Expand is true, the record will be expanded, otherwise it will be collapsed. If EngineID is 0, then all rows encompassed by specified synthesized RecordID will be either expanded or collapsed.

**Update(BSTR i_EngineURL,
GrooveCollectionUpdateOp
i_Operation, void *
i_pUpdateRecord,
IGrooveElement *
io_pUpdateContext);**

Updates the collection. i_Operation is one of : OP_ADD, OP_DELETE, or OP_UPDATE.

**UseSort(BSTR i_SortName,
VARIANT_BOOL**

Sets the sort order for the collection to the named sort order. The specified SortName must be one

TABLE 40

interface IGrooveCollectionCursor : IGrooveDispatch

Move(GrooveCollectionCursorPosition
i_AbsolutePosition,
GrooveCollectionNavigationOp
i_Navigator, long i_Distance, long
*o_pDistanceMoved);

Moves the cursor in either an absolute
or relative amount.

AbsolutePosition may have the values
First, Last, or Current.

Navigator may have the following
values:

Value

Description

NextAny, PriorAny

Move the cursor to the next/previous
source row, traversing down through
child rows and up through parent rows.

NextPeer, PriorPeer

Move the cursor to the next/previous
source row at the same level, stopping
if a row at a higher level is reached.

NextParent, PriorParent

Move the cursor to the next/previous
parent source row, traversing until the
root row is reached.

NextData, PriorData

Move the cursor to the next/previous row that contains a data record.

NextUnread, PriorUnread

Move the cursor to the next/previous unread row.

Distance sets the numbers of iterations to move the cursor, starting at AbsolutePosition and moving through Distance iterations of Navigator movement.

Move returns the number of iterations the cursor was actually moved.

OpenRecord (IGrooveRecord **
o_ppRecord);

Returns an interface pointer to the record the cursor is currently set at.

The following tables illustrate allowed values for the enumerated data types listed in the above interfaces. In particular, Table 41, illustrates allowed values for the GrooveCollationOrder enumerated data type:

TABLE 41

GrooveCollationOrder

CollateAscending	Ordered by ascending data values.
CollateDescending	Ordered by descending data values.
CollateOrdinal	Ordered by ordinal position.

Table 42 illustrates the allowed values for the GrooveCollectionNavigationOp enumerated data type:

TABLE 42

GrooveCollectionNavigationOp

NextAny	Move the cursor to the next source row, traversing down through child rows and up through parent rows.
PriorAny	Move the cursor to the previous source row, traversing down through child rows and up through parent rows.
NextPeer	Move the cursor to the next source row at the same level, stopping if a row at a higher level is reached.
PriorPeer	Move the cursor to the previous source row at the same level, stopping if a row at a higher level is reached.
NextParent	Move the cursor to the next parent source row, traversing until the root row is reached.
PriorParent	Move the cursor to the previous parent source row, traversing until the root row is reached.
NextData	Move the cursor to the next row that contains a data record.

PriorData	Move the cursor to the previous row that contains a data record.
NextUnread	Move the cursor to the next unread row.
PriorUnread	Move the cursor to the next unread row.

Table 43 illustrates the allowed values for the GrooveCollectionCursorPosition enumerated data type:

5

TABLE 43

GrooveCollectionCursorPosition	
First	The first row in the collection.
Last	The last row in the collection.
Current	The current row in the collection. This position is useful for performing relative cursor movement.

Table 44 illustrates the allowed values for the GrooveCollectionRowType enumerated data type:

10

TABLE 44

GrooveCollectionRowType	
ROW_DATA	A row with data values.
ROW_HEADER	A row header, for example, column break values.
ROW_FOOTER	A row footer, for example, column break

values and an aggregated result.

Table 45 illustrates the allowed values for the GrooveCollectionSynthType enumerated data type:

5

TABLE 45

GrooveCollectionSynthType

BreakUnique	Synthesized collection row indicates a change in value of categorized or sorted column. One of the other columns will have the new value.
BreakUnitDay	Synthesized collection row is a break on the change in units of days.
BreakUnitWeek	Synthesized collection row is a break on the change in units of weeks.
BreakUnitMonth	Synthesized collection row is a break on the change in units of months.
BreakUnitYear	Synthesized collection row is a break on the change in units of years.
FuncTotal	Synthesized collection row is the result of an aggregate total function.
FuncCount	Synthesized collection row is the result of an aggregate count function.

Table 46 illustrates the allowed values for the GrooveCollectionUpdateOp enumerated data type:

TABLE 46

GrooveCollectionUpdateOp

OP_ADD	Add the record to the collection.
OP_DELETE	Delete the record from the collection.
OP_UPDATE	Change values of specific fields in this record, which is already in the collection.
OP_REPARENT	Change this record's parent.
OP_CHANGE_ORDINAL	Change the ordinal position of this record in the collection.

Table 47 illustrates the allowed values for the GrooveCollectionWaffleSystem enumerated data type:

TABLE 47

GrooveCollectionWaffleSystemColumns

WAFFLE_ROWTYPE_COLUMN	One of the values for GrooveCollectionRowType.
WAFFLE_SYNTHKIND_COLUMN	If not a data row, one of the values in GrooveCollectionSynthType.

TABLE 48

GrooveCollectionRecordID

NULL_RECORD_ID	The reserved value for the special null record id.
----------------	--

Table 49 illustrates the allowed values for the GrooveSortOrder enumerated data type:

TABLE 49

GrooveSortOrder

Ascending	Collate by ascending data values
Descending	Collate by descending data values.

A software implementation of the above-described embodiment may comprise a series of computer instructions either fixed on a tangible medium, such as a computer readable media, e.g. a diskette, a CD-ROM, a ROM memory, or a fixed disk, or transmissible to a computer system, via a modem or other interface device over a medium. The medium can be either a tangible medium, including, but not limited to, optical or analog communications lines, or may be implemented with wireless techniques, including but not limited to microwave, infrared or other transmission techniques. It may also be the Internet. The series of computer instructions embodies all or part of the functionality previously described herein with respect to the invention. Those skilled in the art will appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Further, such instructions may be stored using any memory technology, present or future, including, but not limited to, semiconductor, magnetic, optical or other memory devices, or transmitted using any communications technology, present or future, including but not limited to optical, infrared, microwave, or other transmission technologies. It is contemplated that such a computer program product

may be distributed as a removable media with accompanying printed or electronic documentation, e.g., shrink wrapped software, pre-loaded with a computer system, e.g., on system ROM or fixed disk, or distributed from a server or electronic bulletin board over a network, e.g., the Internet or World Wide Web.

5 Although an exemplary embodiment of the invention has been disclosed, it will be apparent to those skilled in the art that various changes and modifications can be made which will achieve some of the advantages of the invention without departing from the spirit and scope of the invention. For example, it will be obvious to those reasonably skilled in the art that, although the description was directed to a particular hardware
10 system and operating system, other hardware and operating system software could be used in the same manner as that described. Other aspects, such as the specific instructions utilized to achieve a particular function, as well as other modifications to the inventive concept are intended to be covered by the appended claims.

What is claimed is:

CLAIMS

- 1 1. Apparatus for representing and managing an XML-compliant document in a
2 memory, the XML-compliant document being composed of a plurality of elements
3 arranged in a nested relationship, the apparatus comprising:
4 a data document including a plurality of element objects, each element
5 object representing a part of the XML-compliant document; and
6 a mechanism for arranging the plurality of element objects in a hierarchy
7 representative of the nested relationship of the elements.
- 1 2. Apparatus as recited in claim 1 wherein at least some of the elements contain
2 textual content and wherein element objects representing the elements contain
3 the textual content.
- 1 3. Apparatus as recited in claim 1 wherein at least some of the elements contain
2 attributes having values and wherein element objects representing the elements
3 contain the attribute values.
- 1 4. Apparatus as recited in claim 3 wherein the attribute values contained in the at
2 least some elements are typed.
- 1 5. Apparatus as recited in claim 3 further comprising an attribute index containing
2 consistent pointers to all element objects containing attribute values.
- 1 6. Apparatus as recited in claim 1 wherein the arranging mechanism comprises
2 database pointers and wherein a database pointer in a parent element object
3 points to child objects of the parent element object in order to arrange the parent
4 object and child objects in a hierarchical relationship.

1 7. Apparatus as recited in claim 1 further comprising a schema document
2 referenced by the data document, the schema document containing content that
3 describes the pattern of element objects and attributes, the existence and
4 structure of document indicies, and commonly used strings in the data document.

1 8. Apparatus as recited in claim 7 wherein the schema document is referenced by
2 an XML processing statement in the data document.

1 9. Apparatus as recited in claim 1 further comprising a binary document object for
2 representing a data document containing binary data.

1 10. Apparatus as recited in claim 1 further comprising a document object for
2 representing the data document.

1 11. Apparatus as recited in claim 10 wherein the document object contains links to
2 other document objects so that the other document objects are sub-documents of
3 the document object.

1 12. Apparatus as recited in claim 1 wherein each of the element objects exports a
2 uniform interface containing methods for manipulating each of the element
3 objects.

1 13. Apparatus for binding program code to portions of an XML-compliant document
2 composed of a plurality of elements, each of which is identified by a tag, the
3 elements being arranged in a nested relationship, the apparatus comprising:
4 a data document including a plurality of element objects, each element
5 object representing a part of the XML-compliant document, the plurality of
6 element objects being arranged in a hierarchy representative of the nested
7 relationship of the elements;

8 a schema document referenced by the data document, the schema
9 document containing a registry which maps a tag identifying one of the elements
10 to a program ID code; and

11 a mechanism that uses the program ID code to construct an object
12 containing the program code.

1 14. Apparatus as recited in claim 13 wherein the registry is a two-column table that
2 maps element tags to program ID codes.

1 15. Apparatus as recited in claim 13 wherein the mechanism is responsive to a
2 method call for retrieving the program ID code for constructing the object
3 containing the program code.

1 16. Apparatus as recited in claim 13 wherein the mechanism is the COM object
2 manager and the program ID code is a ProgID code.

1 17. Apparatus as recited in claim 16 wherein the COM manager comprises a locating
2 mechanism that uses the ProgID code to locate the program code and an object
3 constructor that constructs an object incorporating the located program code.

1 18. Apparatus as recited in claim 13 wherein the schema document is referenced in
2 the data document by an XML processing statement.

1 19. Apparatus for representing and managing an XML-compliant document in a
2 memory, the XML-compliant document being updated concurrently by a first
3 process having a first address space in the memory and second process having
4 a second address space in the memory, the apparatus comprising:
5 a first storage manager controlled by the first process that constructs, from
6 class code in the first address space, at least one document object including first

7 data representing a part of the XML-compliant document stored in the first
8 address space;
9 a second storage manager controlled by the second process that
10 constructs, from class code in the second address space which class code is
11 identical to the class code in the first address space, at least one document
12 object including second data representing a part of the XML-compliant document
13 stored in the second address space; and
14 a synchronization mechanism that insures that the first data and the
15 second data are continually equated.

1 20. Apparatus as recited in claim 19 wherein the first data is stored in a region
2 mapped into the first address space and the second data is stored in the same
3 region mapped into the second address space and the synchronization
4 mechanism continually equates the region data mapped in the first and second
5 address spaces.

1 21. Apparatus as recited in claim 20 wherein the second process comprises a
2 mechanism for requesting a copy of the region data from the first address space
3 if the second address space does not have the most recent copy of the region
4 data.

1 22. Apparatus as recited in claim 20 wherein the first process comprises methods for
2 requesting that the synchronization manager lock the region data when the first
3 process is changing the region data in the first address space.

1 23. Apparatus as recited in claim 20 wherein the second process comprises methods
2 for requesting that the synchronization manager lock the region data when the
3 second process is changing the region data in the second address space.

1 24. Apparatus as recited in claim 20 wherein the first process can perform read and
2 write operations on the region and wherein the apparatus further comprises a
3 mechanism for grouping a plurality of the read and write operations into a
4 transaction.

1 25. Apparatus as recited in claim 24 wherein the first process comprises methods for
2 requesting that the synchronization manager lock the region data during the
3 processing of all read and write operations in a transaction.

1 26. Apparatus as recited in claim 25 further comprising a logging system that
2 periodically writes recovery log entries to a persistent database during the
3 processing of all read and write operations in a transaction.

1 27. Apparatus as recited in claim 19 wherein the first process comprises a storage
2 mechanism for storing a copy of the region data in a non-volatile store.

1 28. Apparatus as recited in claim 27 wherein the non-volatile store comprises an
2 object store.

1 29. Apparatus as recited in claim 27 wherein the non-volatile store comprises a file
2 system.

1 30. Apparatus as recited in claim 19 wherein the synchronization mechanism
2 comprises a distributed memory system.

1 31. Apparatus as recited in claim 19 wherein both the first and second address
2 spaces contain equivalent program code for manipulating the first and second
3 document objects.

1 32. Apparatus as recited in claim 19 wherein the first and second storage manager
2 each construct a cross-process synchronization object that is used to
3 synchronize the first and second processes.

1 33. Apparatus for representing and managing an XML-compliant document in a
2 memory, the XML-compliant document being composed of a plurality of
3 elements, the elements being arranged in a nested relationship, the apparatus
4 comprising:

5 a data document including a plurality of element objects, each element
6 object representing a part of the XML-compliant document; and

7 a collection manager that maps the element objects into a tabular data
8 structure including an index structure.

1 34. Apparatus as recited in claim 33 further comprising a record set engine that is
2 responsive to user commands for propagating a set of updates for the tabular
3 data structure to the collection manager.

1 35. Apparatus as recited in claim 34 wherein the collection manager further
2 comprises an update mechanism which responds the set of updates by updating
3 the index structure.

1 36. Apparatus as recited in claim 35 wherein the collection manager further
2 comprises a notification system that notifies the users when changes are made
3 to the tabular data structure.

1 37. Apparatus as recited in claim 36 wherein the collection manager further
2 comprises a navigation mechanism for creating a cursor to allow the users to
3 navigate within the tabular data structure.

1 38. A method for representing and managing an XML-compliant document in a
2 memory, the XML-compliant document being composed of a plurality of elements
3 arranged in a nested relationship, the method comprising:

4 (a) creating a data document in the memory including a plurality of element
5 objects, each element object representing a part of the XML-compliant
6 document; and

7 (b) arranging the plurality of element objects in a hierarchy representative of
8 the nested relationship of the elements.

1 39. A method as recited in claim 38 wherein at least some of the elements contain
2 textual content and wherein element objects representing the elements contain
3 the textual content.

1 40. A method as recited in claim 38 wherein at least some of the elements contain
2 attributes having values and wherein element objects representing the elements
3 contain the attribute values.

1 41. A method as recited in claim 40 wherein the attribute values contained in the at
2 least some elements are typed.

1 42. A method as recited in claim 40 further comprising an attribute index containing
2 consistent pointers to all element objects containing attribute values.

1 43. A method as recited in claim 38 wherein step (b) comprises creating a database
2 pointer in a parent element object which pointer points to child objects of the
3 parent element object in order to arrange the parent object and child objects in a
4 hierarchical relationship.

1 44. A method as recited in claim 38 further comprising (c) creating a schema
2 document referenced by the data document in the memory, the schema
3 document containing content that describes the pattern of element objects and
4 attributes, the existence and structure of document indicies, and commonly used
5 strings in the data document.

1 45. A method as recited in claim 44 wherein step (c) comprises creating the schema
2 document referenced by an XML processing statement in the data document.

1 46. A method as recited in claim 38 further comprising (d) creating a binary
2 document object in the memory for representing a data document containing
3 binary data.

1 47. A method as recited in claim 38 further comprising (e) creating a document
2 object in the memory for representing the data document.

1 48. A method as recited in claim 47 wherein the document object contains links to
2 other document objects so that the other document objects are sub-documents of
3 the document object.

1 49. A method as recited in claim 38 wherein each of the element objects exports a
2 uniform interface containing methods for manipulating each of the element
3 objects.

1 50. A method for binding program code in a memory to portions of an XML-compliant
2 document composed of a plurality of elements, each of which is identified by a
3 tag, the elements being arranged in a nested relationship, the XML-compliant
4 document being stored in the memory and the method comprising:
5 (a) creating a data document in the memory including a plurality of element
6 objects, each element object representing a part of the XML-compliant

document, the plurality of element objects being arranged in a hierarchy representative of the nested relationship of the elements;

(b) creating a schema document referenced by the data document in the memory, the schema document containing a registry which maps a tag identifying one of the elements to a program ID code; and

(c) using the program ID code to construct an object containing the program code.

51. A method as recited in claim 50 wherein the registry is a two-column table that maps element tags to program ID codes.

52. A method as recited in claim 50 wherein step (c) is initiated in response to a method call for retrieving the program ID code.

53. A method as recited in claim 50 wherein step (c) is performed by the COM object manager and the program ID code is a ProglD code.

54. A method as recited in claim 53 wherein the COM manager performs the steps of using the ProglD code to locate the program code and constructing an object incorporating the located program code.

55. A method as recited in claim 50 wherein the schema document is referenced in the data document by an XML processing statement.

56. A method for representing and managing an XML-compliant document in a memory, the XML-compliant document being updated concurrently by a first process having a first address space in the memory and second process having a second address space in the memory, the method comprising:

(a) using a first storage manager controlled by the first process to construct, from class code in the first address space, at least one document object

- including first data representing a part of the XML-compliant document stored in the first address space;
- (b) using a second storage manager controlled by the second process to construct, from class code in the second address space which class code is identical to the class code in the first address space, at least one document object including second data representing a part of the XML-compliant document stored in the second address space; and
- (c) insuring that the first data and the second data are continually equated.

57. A method as recited in claim 56 wherein the first data is stored in a region mapped into the first address space and the second data is stored in the same region mapped into the second address space and step (c) comprises continually equating the region data mapped in the first and second address spaces.

58. A method as recited in claim 57 wherein step (c) comprises requesting a copy of the region data from the first address space if the second address space does not have the most recent copy of the region data.

59. A method as recited in claim 57 wherein step (c) comprises locking the region data when the first process is changing the region data in the first address space.

60. A method as recited in claim 57 wherein step (c) comprises locking the region data when the second process is changing the region data in the second address space.

~~Apparatus~~ A method as recited in claim 57 wherein the first process can perform read and write operations on the region and wherein the method further comprises (d) grouping a plurality of the read and write operations into a transaction.

KGM
ZLR
EJP
R.O.

- 1 62. A method as recited in claim 61 wherein step (c) comprises locking the region
2 data during the processing of all read and write operations in a transaction.
- 1 63. A method as recited in claim 62 wherein step (c) further comprises periodically
2 writing recovery log entries to a persistent database during the processing of all
3 read and write operations in a transaction.
- 1 64. A method as recited in claim 56 further comprising (e) under the control of the
2 first process, storing a copy of the region data in a non-volatile store.
- 1 65. A method as recited in claim 64 wherein the non-volatile store comprises an
2 object store.
- 1 66. A method as recited in claim 64 wherein the non-volatile store comprises a file
2 system.
- 1 67. A method as recited in claim 56 wherein step (c) is performed by a distributed
2 memory system.
- 1 68. A method as recited in claim 56 further comprising (f) manipulating the first and
2 second document objects with equivalent program code in both the first and
3 second address spaces.
- 1 69. A method as recited in claim 56 further comprising (g) constructing a cross-
2 process synchronization object that is used to synchronize the first and second
3 processes.
- 1 70. A method for representing and managing an XML-compliant document in a
2 memory, the XML-compliant document being composed of a plurality of

elements, the elements being arranged in a nested relationship, the method comprising:

- (a) creating a data document in the memory including a plurality of element objects, each element object representing a part of the XML-compliant document; and
- (b) mapping the element objects into a tabular data structure including an index structure.

71. A method as recited in claim 70 further comprising (c) propagating a set of updates for the tabular data structure to the collection manager in response to user commands.

72. A method as recited in claim 71 further comprising (d) updating the index structure in response to set of updates.

73. A method as recited in claim 72 further comprising (e) notifying the users when changes are made to the tabular data structure.

74. A method as recited in claim 73 further comprising (f) creating a cursor to allow the users to navigate within the tabular data structure.

75. A computer program product for representing and managing an XML-compliant document in a memory, the XML-compliant document being composed of a plurality of elements arranged in a nested relationship, the computer program product comprising a computer usable medium having computer readable program code thereon, including:

- program code for creating a data document in the memory including a plurality of element objects, each element object representing a part of the XML-compliant document; and

9 program code for arranging the plurality of element objects in a hierarchy
10 representative of the nested relationship of the elements.

1 76. A computer program product for binding program code in a memory to portions of
2 an XML-compliant document composed of a plurality of elements, each of which
3 is identified by a tag, the elements being arranged in a nested relationship, the
4 XML-compliant document being stored in the memory and the computer program
5 product comprising a computer usable medium having computer readable
6 program code thereon, including:

7 program code for creating a data document in the memory including a
8 plurality of element objects, each element object representing a part of the XML-
9 compliant document, the plurality of element objects being arranged in a
10 hierarchy representative of the nested relationship of the elements;

11 program code for creating a schema document referenced by the data
12 document in the memory, the schema document containing a registry which
13 maps a tag identifying one of the elements to a program ID code; and

14 program code for using the program ID code to construct an object
15 containing the program code.

1 77. A computer program product for representing and managing an XML-compliant
2 document in a memory, the XML-compliant document being updated
3 concurrently by a first process having a first address space in the memory and
4 second process having a second address space in the memory, the computer
5 program product comprising a computer usable medium having computer
6 readable program code thereon, including:

7 program code for using a first storage manager controlled by the first
8 process to construct, from class code in the first address space, at least one
9 document object including first data representing a part of the XML-compliant
10 document stored in the first address space;

11 program code for using a second storage manager controlled by the
12 second process to construct, from class code in the second address space which
13 class code is identical to the class code in the first address space, at least one
14 document object including second data representing a part of the XML-compliant
15 document stored in the second address space; and
16 program code for insuring that the first data and the second data are
17 continually equated.

1 78. A computer program product for representing and managing an XML-compliant
2 document in a memory, the XML-compliant document being composed of a
3 plurality of elements, the elements being arranged in a nested relationship, the
4 computer program product comprising a computer usable medium having
5 computer readable program code thereon, including:

6 program code for creating a data document in the memory including a
7 plurality of element objects, each element object representing a part of the XML-
8 compliant document; and

9 program code for mapping the element objects into a tabular data
10 structure including an index structure.

1 79. A computer data signal embodied in a carrier wave for representing and
2 managing an XML-compliant document in a memory, the XML-compliant
3 document being composed of a plurality of elements arranged in a nested
4 relationship, the computer data signal comprising:

5 program code for creating a data document in the memory including a
6 plurality of element objects, each element object representing a part of the XML-
7 compliant document; and

8 program code for arranging the plurality of element objects in a hierarchy
9 representative of the nested relationship of the elements.

80. A computer data signal embodied in a carrier wave for binding program code in a memory to portions of an XML-compliant document composed of a plurality of elements, each of which is identified by a tag, the elements being arranged in a nested relationship, the XML-compliant document being stored in the memory and the computer data signal comprising:

program code for creating a data document in the memory including a plurality of element objects, each element object representing a part of the XML-compliant document, the plurality of element objects being arranged in a hierarchy representative of the nested relationship of the elements;

program code for creating a schema document referenced by the data document in the memory, the schema document containing a registry which maps a tag identifying one of the elements to a program ID code; and

program code for using the program ID code to construct an object containing the program code.

81. A computer data signal embodied in a carrier wave for representing and managing an XML-compliant document in a memory, the XML-compliant document being updated concurrently by a first process having a first address space in the memory and second process having a second address space in the memory, the computer data signal comprising:

program code for using a first storage manager controlled by the first process to construct, from class code in the first address space, at least one document object including first data representing a part of the XML-compliant document stored in the first address space;


program code for using a second storage manager controlled by the second process to construct, from class code in the second address space which class code is identical to the class code in the first address space, at least one document object including second data representing a part of the XML-compliant document stored in the second address space; and

15 program code for insuring that the first data and the second data are
16 continually equated.

- 1 82. A computer data signal for representing and managing an XML-compliant
2 document in a memory, the XML-compliant document being composed of a
3 plurality of elements, the elements being arranged in a nested relationship, the
4 computer data signal comprising:
5 program code for creating a data document in the memory including a
6 plurality of element objects, each element object representing a part of the XML-
7 compliant document; and
8 program code for mapping the element objects into a tabular data
9 structure including an index structure.

ABSTRACT OF THE DISCLOSURE

An in-memory storage manager represents XML-compliant documents as a collection of objects in memory. The collection of objects allows the storage manager to manipulate the document, or parts of the document with a consistent interface and to provide for features that are not available in conventional XML documents, such as element attributes with types other than text and documents that contain binary rather than text information. In addition, in the storage manager, the XML-compliant document is associated with a schema document which defines the arrangement of the document elements and attributes. The schema data associated with a document can contain a mapping between document elements and program code to be associated with each element. The storage manager further has methods for retrieving the code from the element tag. The retrieved code can then be invoked using attributes and content from the associated element and the element then acts like a conventional object. Further, the storage manager allows real-time access by separate process operating in different contexts. The objects that are used to represent the document are constructed from common code found locally in each process. In addition, the data in the objects is also stored in memory local to each process. The local memories are synchronized by means of a distributed memory system that continually equates the data copies of the same element in different processes. Client-specified collections are managed by a separate collection manager. The collection manager maintains a data structure called a "waffle" that represents the XML data structures in tabular form. A record set engine that is driven by user commands propagates a set of updates for a collection to the collection manager. Based on those updates, the collection manager updates index structures and may notify waffle users via the notification system.






FIG. 2




FIG. 3

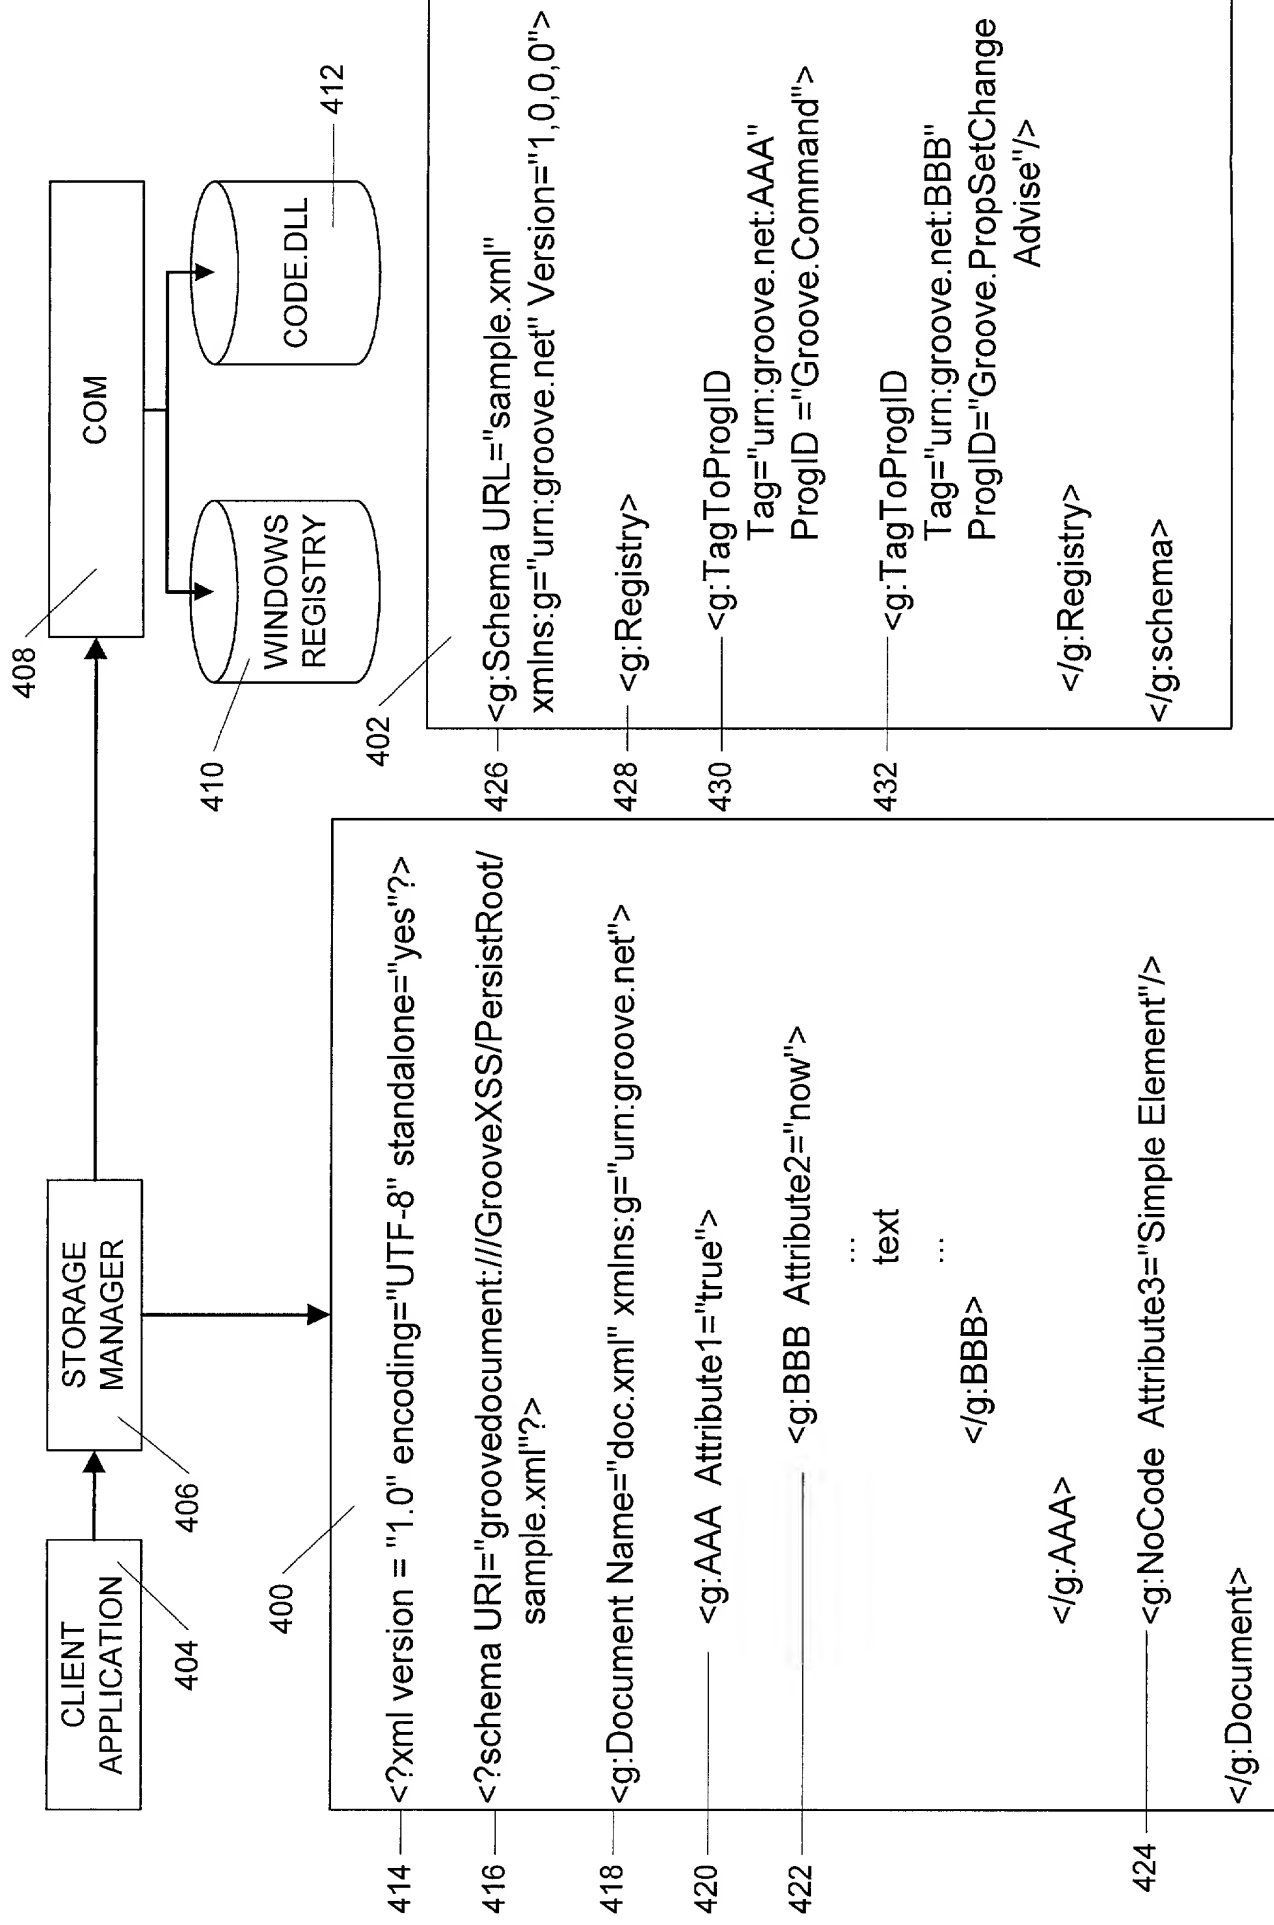


FIG. 4A

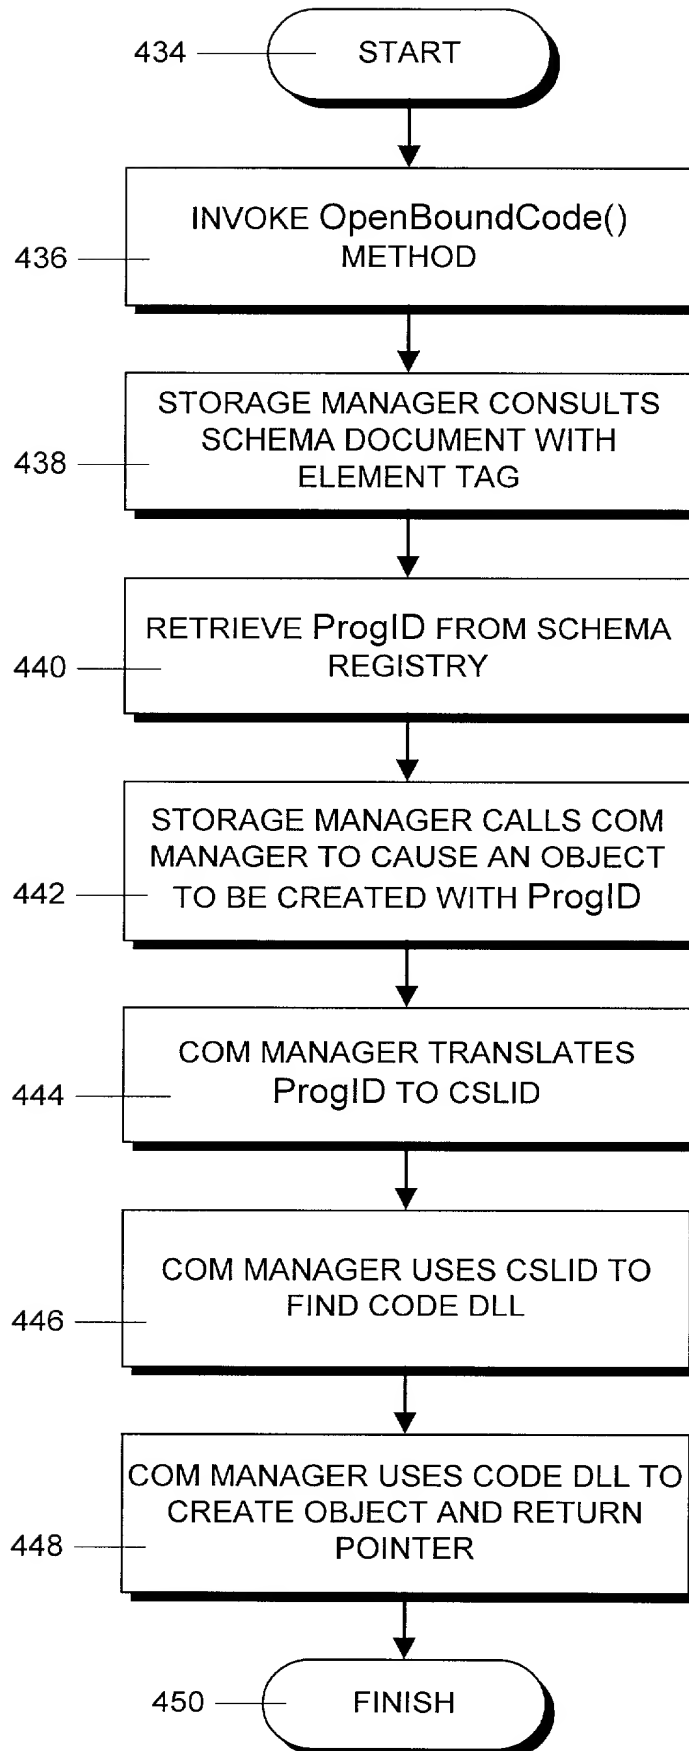


FIG. 4B

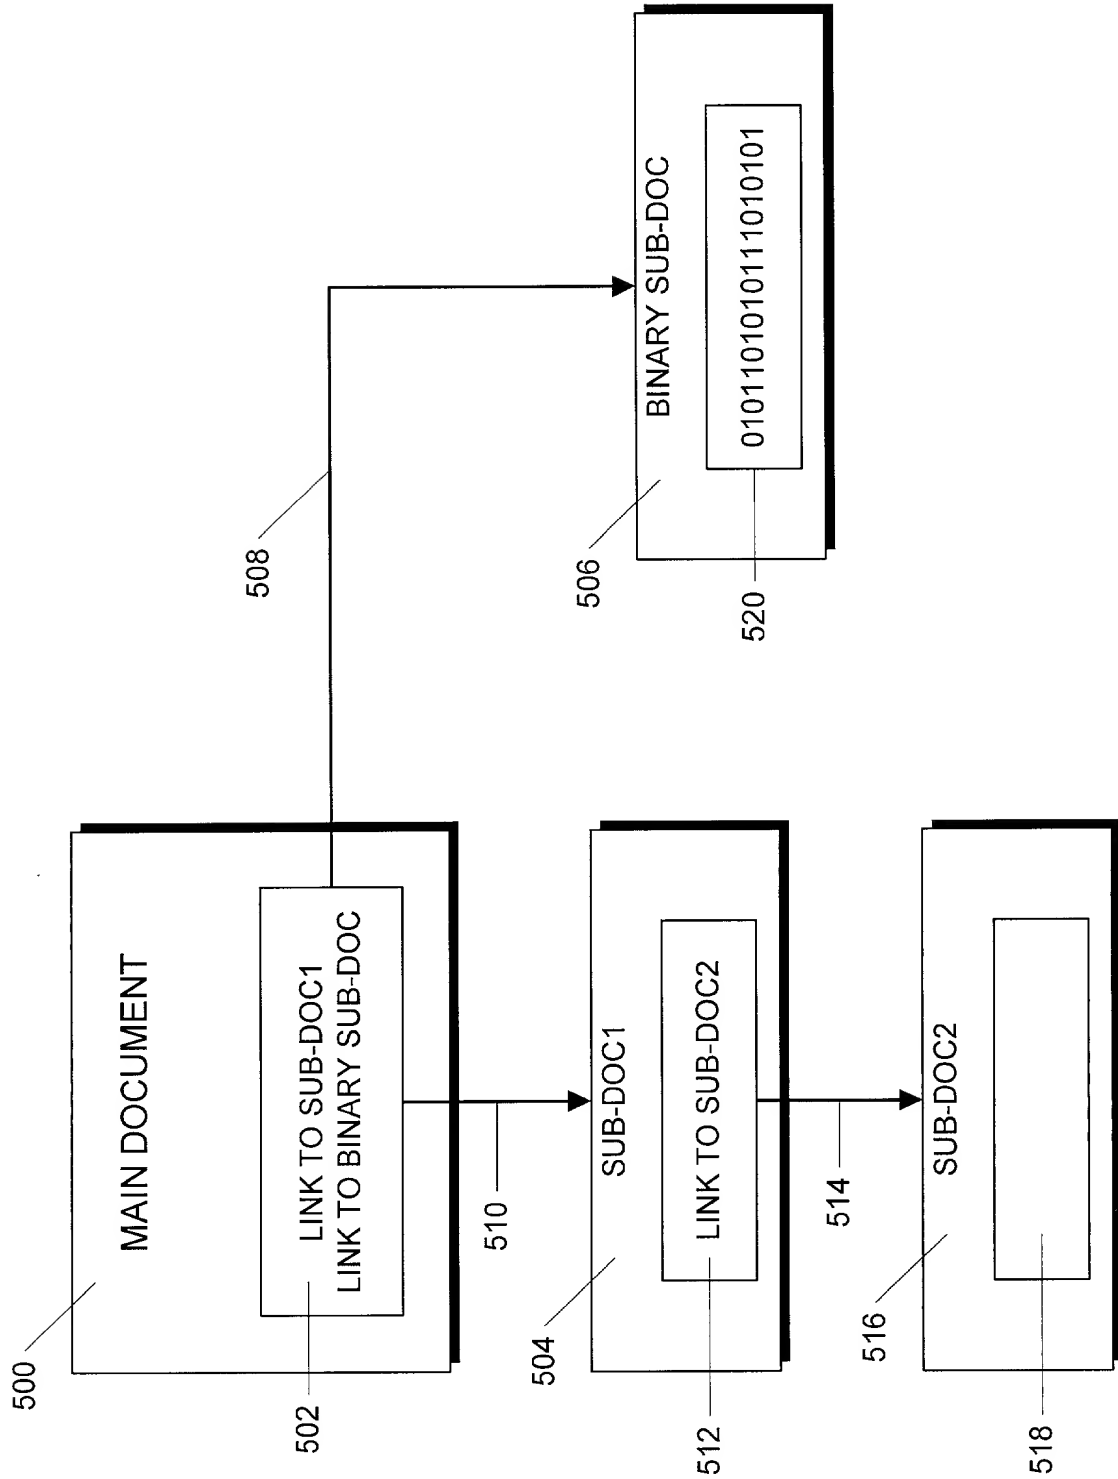


FIG. 5




FIG. 6




FIG. 7

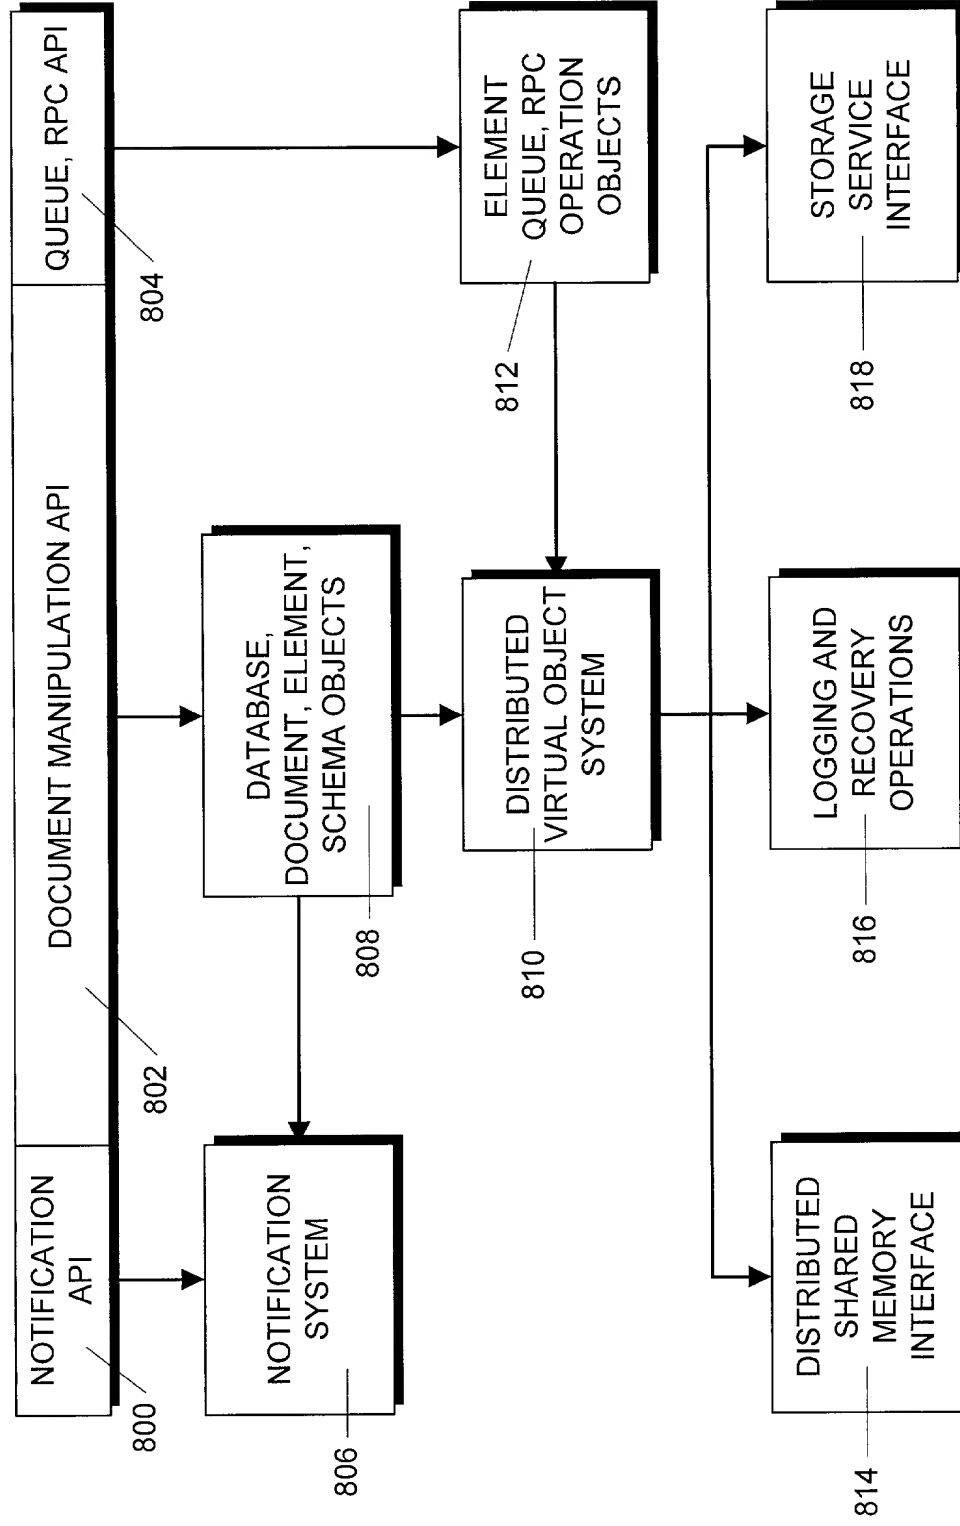


FIG. 8




FIG. 9

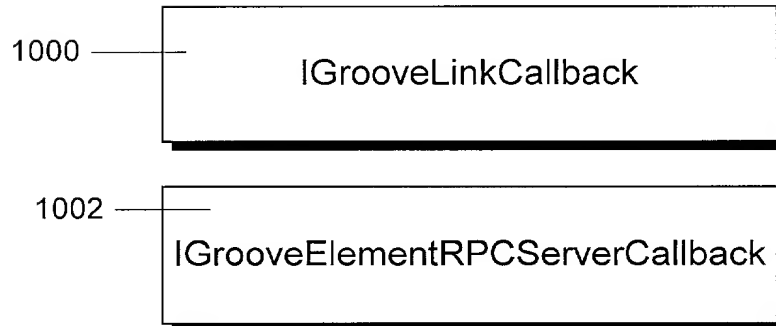


FIG. 10

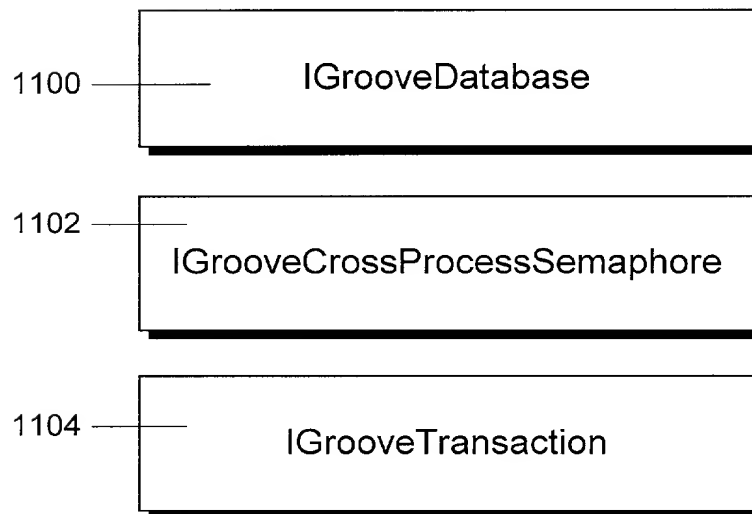


FIG. 11

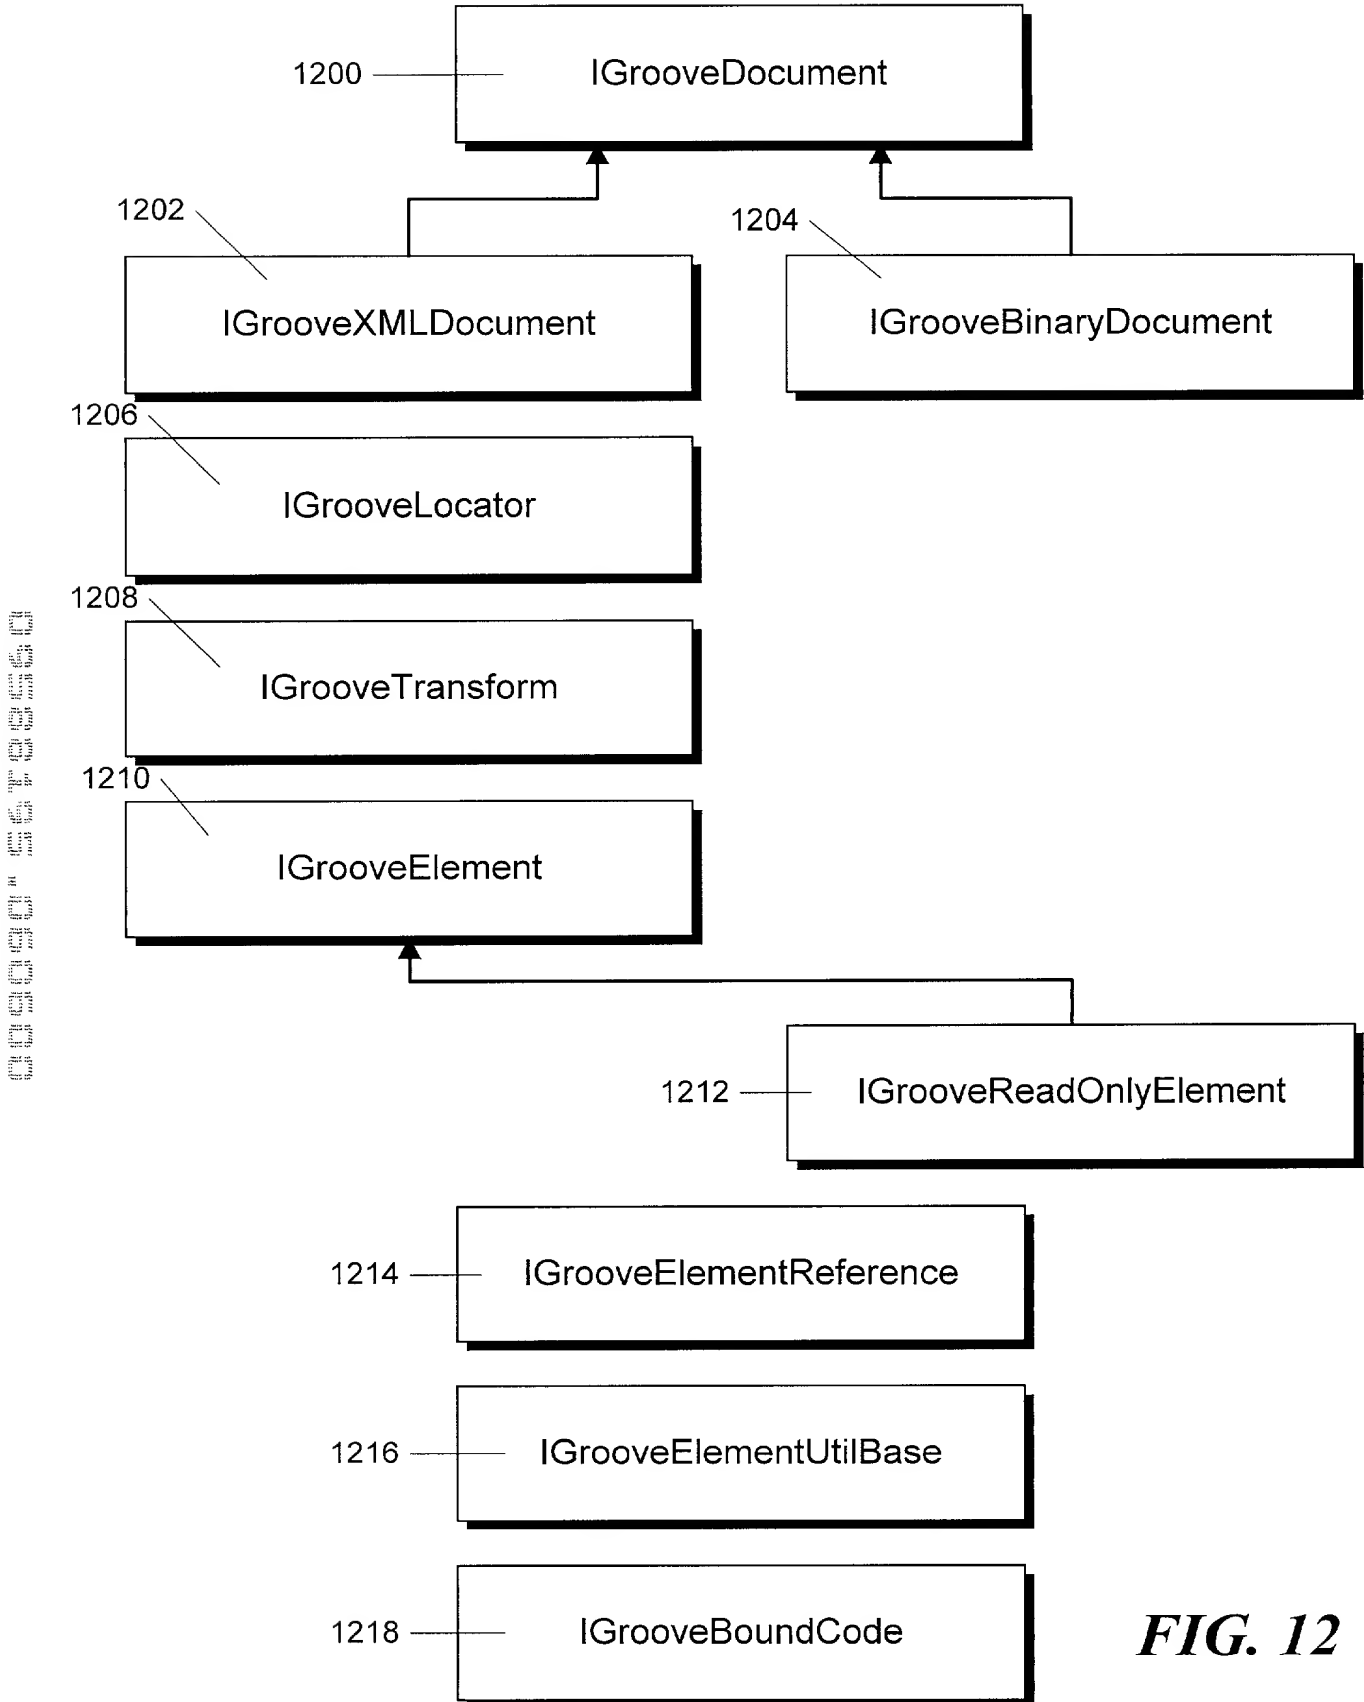


FIG. 12

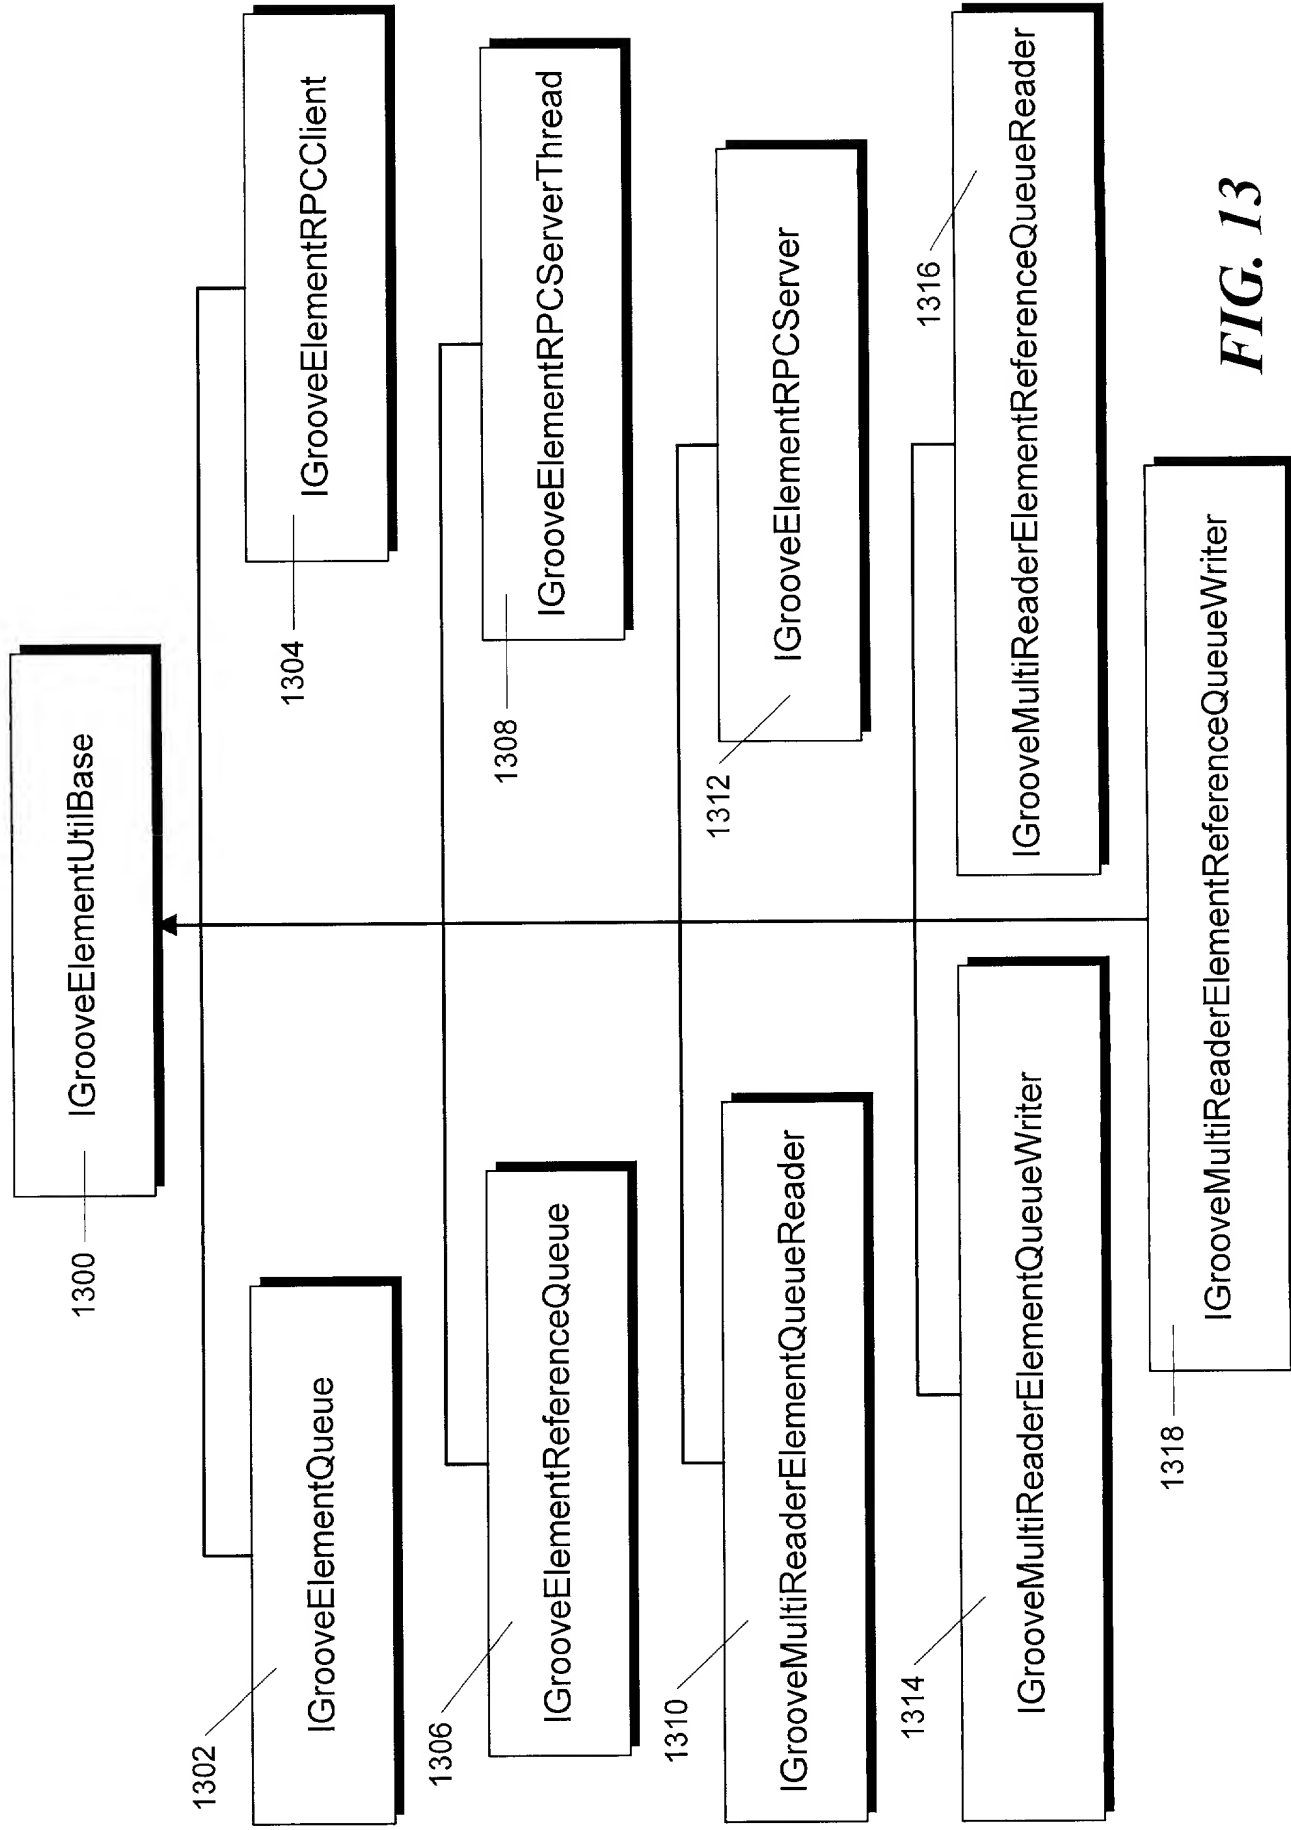


FIG. 13

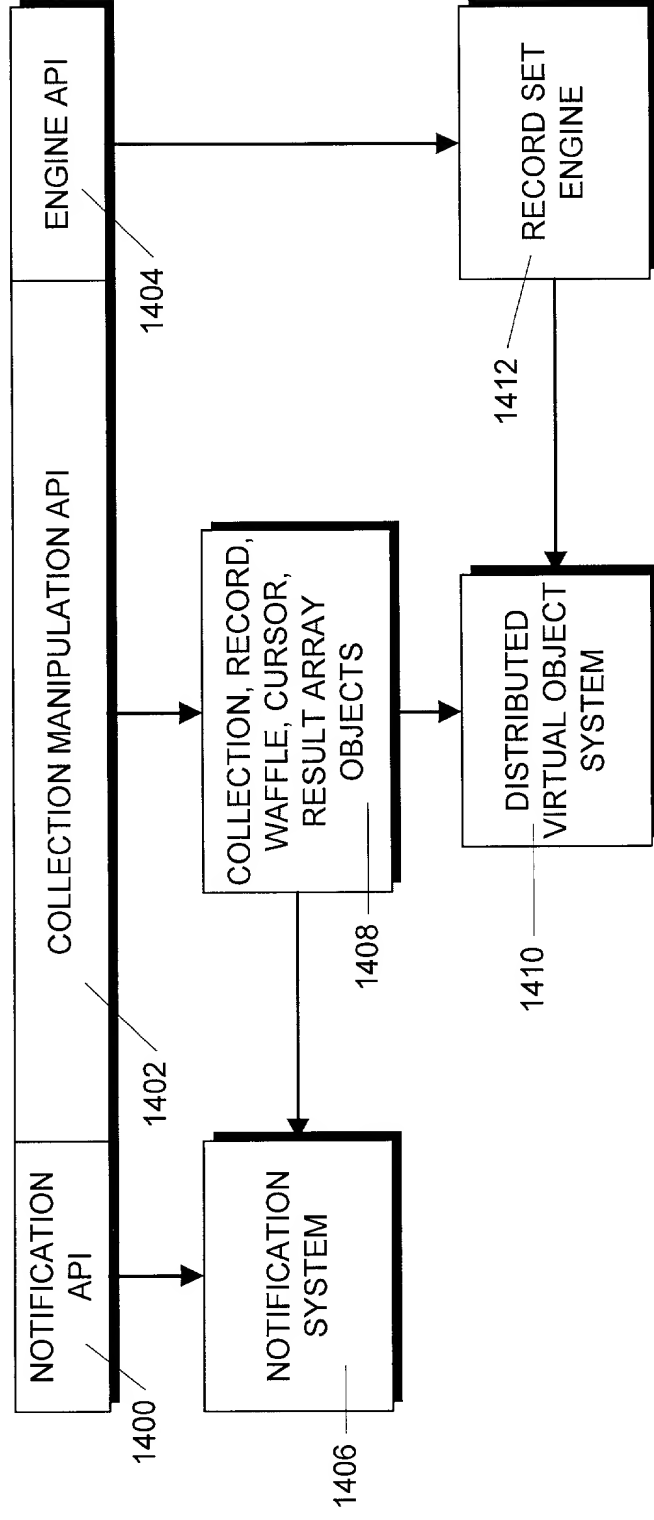


FIG. 14




FIG. 15

DECLARATION AND POWER OF ATTORNEYDocket No. **G0008/7003**

As a below-named inventor, I hereby declare that:

1. My residence, post-office address and citizenship are as stated below next to my name.
2. I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled, **METHOD AND APPARATUS FOR EFFICIENT MANAGEMENT OF XML DOCUMENTS**, the specification of which is attached hereto and identified by Docket No. G0008/7003.
3. I have reviewed and understand the contents of the above-identified application specification, including the claims.
4. I acknowledge the duty to disclose all information known to me that is material to patentability as defined in 37 C.F.R. §1.56.

-
5. I hereby claim foreign priority benefits under 35 U.S.C. §119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate or 365(a) of any PCT application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the appropriate box, any foreign application for patent or inventor's certificate, or any PCT international application having a filing date before that of the application on which priority is claimed:

Application No.	Country	Filing Date	Priority NOT Claimed	Certified Copy Attached.
			<input type="checkbox"/>	<input type="checkbox"/>

☐ Additional foreign application numbers are listed on a supplemental priority data sheet attached hereto

-
6. I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional applications listed below:

Application No.	Filing Date
-----------------	-------------

☐ Additional provisional application numbers are listed on a supplemental data sheet attached hereto

-
7. I hereby claim the benefit under 35 U.S.C. §120, of the United States Application(s) or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. §112, I acknowledge the duty to disclose all information which is material to patentability as defined in 37 C.F.R. §1.56, and which became available to me between the filing date of the prior application and the national or PCT international filing date of this application:

Application No.	Filing Date	Parent Patent No.
-----------------	-------------	-------------------

☐ Additional U.S. or PCT application numbers are listed on a supplemental data sheet attached hereto

8. I hereby appoint the attorneys listed under the KUDIRKA & JOBSE, LLP customer number:

021127



jointly, and each of them severally, its attorneys at law, with full power of substitution, delegation and revocation, to prosecute this application to register, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected therewith. Address all correspondence to

Paul E. Kudirka, Esq.

at the customer address for the customer number listed above and

telephone no. (617) 367-4600; facsimile number (617) 367-4656.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both under 18 U.S.C. §1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

First Inventor Name: Raymond E. Ozzie

Inventor's Signature:

A handwritten signature in dark ink, appearing to be 'Raymond E. Ozzie', written over a horizontal line.

Date:

July 5, 2000

Citizenship:

US

Residence Address:

50 Harbor Street, Manchester, Massachusetts 01944

Post Office Address:

50 Harbor Street, Manchester, Massachusetts 01944

Second Inventor Name: Kenneth G. Moore

Inventor's Signature:

A handwritten signature in dark ink, appearing to be 'Kenneth G. Moore', written over a horizontal line.

Date:

June 5, 2000

Citizenship:

US

Residence Address:

7 Jack Rabbit Lane, Westford, Massachusetts 01886

Post Office Address:

7 Jack Rabbit Lane, Westford, Massachusetts 01886

☒ Additional inventors are being named on the additional inventor sheet attached hereto.

DECLARATION - SUPPLEMENTAL INVENTOR SHEET

Third Inventor Name: Ransom L. Richardson

Inventor's Signature: Norman H. Tiers Date: June 5, 2000

Citizenship: US

Residence Address: 102 Morrison Avenue, Apt. #2, Somerville, Massachusetts 02144

Post Office Address: 102 Morrison Avenue, Apt. #2, Somerville, Massachusetts 02144

Fourth Inventor Name: Edward J. Fischer

Inventor's Signature: Edmund J. Frisch Date: June 5, 2000

Citizenship: US

Residence Address: 177 Pemberton Street, #8, Cambridge, Massachusetts 02140

Post Office Address: 177 Pemberton Street, #8, Cambridge, Massachusetts 02140